TEACHING SIGNAL PROCESSING ONLINE: A REPORT FROM THE TRENCHES

Don H. Johnson, Paolo Prandoni, Pedro C. Pinto, Martin Vetterli

Department of Electrical and Computer Engineering
Rice University, Houston, TX 77005, USA.

School of Computer and Communication Sciences
Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland
dhj@rice.edu, {paolo.prandoni, pedro.pinto, martin.vetterli}@epfl.ch

ABSTRACT

After years of experimentation, online teaching has gone through a phase transition with the appearance of massive open online courses (MOOCs), following the model pioneered by Salman Khan with short videos and the use of tablets. Dozens of university-level courses are now available, and followed by hundreds of thousands of online students. We report on early experiences with teaching "Fundamentals of Electrical Engineering" and "Digital Signal Processing" on an open online platform (Coursera). We address in particular: (i) suitability of online platforms for signal processing oriented topics; (ii) structuring of material for the online format; (iii) quizzes and exercises for large classes; (iv) grading methods and possible certification issues; (v) learning from teaching, forums, and data mining of students feedback.

Index Terms— Online education, signal processing.

1. INTRODUCTION: FROM SOCRATES TO MOOCS

How to transmit knowledge is a challenge as old as civilization itself. Teaching, in its various forms, aims at the core of spreading knowledge. One of the earliest forms of teaching was the one used by Socrates in the School of Athens: the teacher discussed fundamental philosophical questions with his students, arguing the points, and allowing truth to emerge through this dialog. In short, Socrates did not believe in books for teaching purposes. Nonetheless, his favorite pupil, Plato, wrote down the Dialogs, which became the pillars of western thought. But until the (re)invention of printing in the 16th century, diffusion of information through books was painstakingly slow, as manuscripts were copied by hand, one by one. After Gutenberg, diffusion of books exploded, as making hundreds of thousands of copies was easy and economically viable. Books were used for teaching, with reference books and teaching manuals being used in universities, leading finally to the notion of textbooks as we know it today. From the introduction of printing to the 20th century, not much happened as far as technology for information diffusion is concerned. With the electronic media revolution of the 20th century, this changed, but without much impact on teaching, surprisingly. Ex cathedra lectures with hundreds of students in a classroom watching a lecturer writing on a blackboard were the norm. Slideshows and laser pointers have not fundamentally changed the delivery method: a broadcast channel with little feedback, other than exams at the end of the semester. With the introduction of the World Wide Web in 1989, the seeds for a revolution were planted. First, knowledge was available for free, at a scope never seen before, with Wikipedia being a prime example. However, the time for the Web to impact on teaching required more time. In 2006, Salman Khan started teaching mathematics to his cousins over the internet using standard technologies (a tablet and video over the internet) and an attractive pedagogy (examples, visuals, interactivity). This became an incredible success, and the Khan Academy provides today several thousand lectures that have been watched close to 200 million times. A similar concept was used to bring online several university-level lectures, first by Stanford in the Fall of 2011, and then by several startups, leading to massive open online courses (MOOCs) with hundreds of thousands of students registered. MIT, Harvard and other universities have also launched edX—a not-for-profit initiative following the Open Courseware effort launched by MIT ten years ago.

The pedagogical concept as of today is the following: a standard classroom lecture is split into short (10-15 minute) video lectures, which include simple quizzes and automatically-graded exercises. The challenge is to determine how we can effectively teach signal processing using this platform. And, can the availability of finer feedback (a granularity in questions, quizzes and exercises that is not practical in the classroom) improve our delivery and transmission of material? In some sense, can we get closer to the Socratic method by using state of the art technology? We will try to answer this question by doing two case studies namely: (i) Fundamentals of Electrical Engineering and (ii) Digital Signal Processing.

2. PLATFORM AND INFRASTRUCTURE

At the time of writing, an instructor who wishes to teach a MOOC can choose between three, roughly equivalent, platforms: Coursera, edX, or Class2Go. Other platforms exist, but currently do not benefit from high visibility or wide distribution. The platform connects teachers with students, and supports the whole MOOCs lifecycle: it provides the tools for managing and delivering the course content, and it hosts all video lectures, assignments, and discussion forums. However, the burden of creating, recording, and producing the course materials is placed on the instructor. While the features of the available platforms are comparable, care should be taken in understanding the policies of each platform, namely regarding copyright, cost, exclusivity, and ownership of student data.

To address the course production task, we set up a dedicated recording studio (Fig. 1). The studio configuration is highly depen-

This work was supported by ERC Advanced Grant–Support for Frontier Research–SPARSAM, Nr: 247006.

Fig. 1. The recording studio for the DSP class—a modern version of the classroom. The instructor teaches on the tablet, while an overhead camera records the instructor's hand. Then, the video and audio streams are merged and edited in the workstation.

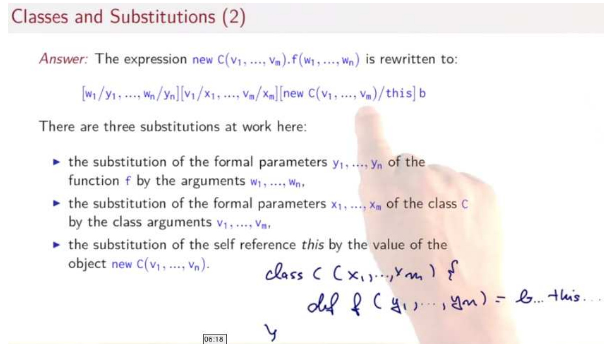

Fig. 2. The video lectures consist of a slideshow, on which the instructor's hand and writing are transparently overlaid.

dent on the desired lecture format. After experimenting with various formats, we decided that the most effective and engaging format was displaying a slide deck with the instructor's hand overlaid semi-transparently (Fig. 2). The instructor can freely point and write on the slides, while explaining the material. We found that showing the instructor's face was distracting and decided not to include it. Our goal was to recreate the experience of sitting next to a colleague and have the material explained to us on a sheet of paper—rather than being lectured in a conventional way. To support such format, the recording studio requires two main components: i) a tablet where the instructor writes on the slides; and ii) an overhead camera that captures the instructor's hand. After each recording session, we simply merge three different streams: the screen recording (tablet), the hand recording (overhead camera), and the instructor's voice (microphone). Substantial production effort is needed to further clean up the video clips, by removing silences and mistakes.

3. TEACHING "FUNDAMENTALS OF ELECTRICAL ENGINEERING" ONLINE

The course titled "Fundamentals of Electrical Engineering" is the first course for majors in Electrical & Computer Engineering at Rice University. The course's structure differs from most introductory courses in that it touches on virtually all topics and applications found in an undergraduate engineering curriculum from a frequency-domain perspective.[1] This course, designed several years ago, focuses on information as a fundamental theme: how do electrical signals represent information, how do systems extract it and how is information communicated with signals [1]. Topics covered include circuit theory, Fourier series and the Fourier transform, analog and digital signal processing, analog and digital communication, and fundamentals of information theory (lossless and lossy compression, error correcting codes and statements of Shannon's theorems). Thus, from an online educational viewpoint, this course would be ideal *for someone with a technical background* to learn about the field of electrical engineering at some technical depth.

Another consideration in offering this course as a MOOC was that the course's text was written for the course and is available for no charge online at Connexions (`http://www.cnx.org/content/col10040/latest/`). Consequently, students could obtain a coherent treatment of electrical engineering, without notation differences and other such annoyances. On the downside, the Rice course, normally taken by sophomore electrical engineering majors and those curious about electrical engineering, is considered to be a very challenging course but not artificially demanding to make it a "weed-out" course. At Rice, a typical week devoted by students to the course consists of three hours of lectures by the instructor, a weekly concept and example session taught by the course's teaching assistant (a graduate student), problem sessions led by course assistants (junior and senior undergraduates) and working problem sets (group work encouraged). The course's pedagogical philosophy is that working problems is key; in some cases, concepts are first introduced in problems.

Mimicking this support network for a MOOC represents a clear challenge. The lectures consist of videos made by the instructor (Johnson). Because of his bad experiences with classroom videos (poor contrast of the white/black board and audio problems), videos were made of slide presentations. Rather than leaving the "talking head" on the screen at all times, the videos focus on slides scribbled on during the lecture. Dynamic annotations on the slides is an important feature: the student knows where the instructor's point-of-focus relates to the slide and writing on the slide slows down the presentation pace.

These videos have also been made available to Rice students taking the course. They find a different viewpoint comes across and helps them understand the material better. However, they do not feel that the videos come anywhere close to the classroom experience. The instructor keys his lectures from student reactions, asks many questions during class and sometimes ventures into territories related to course material but not directly addressed—for example, what is the quality control variable in JPEG compression and how does JPEG compression work in general terms? Such spur-of-the-moment elaborations are not possible with videos made of slide presentations.

The most difficult and trying aspect of creating a MOOC is the homework. Naively, we thought that the homework assigned to Rice students could also be assigned to MOOC students. Creating questions, both in-video quizzes and homework problems, was easy

[1]The time-domain viewpoint along with the concept of stability is addressed in a subsequent course.

in the Coursera platform.[2] However, grading/assessing student answers is not easily accomplished. The easiest question style to ask is a multiple-choice question. While this might suffice for in-video quizzes meant to probe the student's grasp of the lecture without affecting their grade, giving a prospective engineer the answer to a design-style question runs counter to having students think hard about choosing an approach and implementing it. Consequently, we wanted our homework questions to be simple questions with no answer provided. Simple expressions (typed as one would in Matlab or C) and numeric answers can be graded within the Coursera system. However, plots, circuits and derivations cannot. Also, expressions involving functions not pre-defined in the Coursera system, such as the unit step and unit sample, cannot be graded by computer. Furthermore, special notation common in DSP is not accommodated: an answer of $\sin(2\pi t)$ can be graded in the current system, but an answer of $\sin[2\pi n]$ cannot because square brackets are not allowed. Hopefully, these limitations will disappear as the system matures. Hand grading is clearly out of the question because of the "M" in MOOC: 18,000 students are registered for "Fundamentals of Electrical Engineering." Consequently, only questions having simple answers can be asked and computer-graded at the moment, lessening the depth of the educational experience. Currently, Rice does not offer a certificate of completion for the course as the University is considering its policies toward MOOCs. However, student grades on homework questions are collected to monitor their understanding and discover what topics need to be taught better.

Providing collaborative environments for students to learn is part of the Coursera system. Forum discussions on the course's Wiki can help students collaborate and solve problems. Highly rated questions that arise in the discussion forums can be the subject of special videos made by the instructor or teaching assistant. It is this personal aspect of teaching a MOOC that demands resources: instructor time and teaching assistants (they could be undergraduates) are needed to monitor the forums and respond to them. Of course, the forums also provide a mechanism for instructor feedback for modifying videos and refining homework problems. The course is not scheduled to go online until the spring of 2013, which means we have no experience with operating the course. We predict that the time required to care for the course beyond slide and video preparation will be comparable to the on-campus version.

At Rice, teaching a course online has to be viewed as an experiment. However, benefits have already occurred with on-campus students that we are currently evaluating. We are considering running a "flipped classroom" for future versions of the course. Here, on-campus students watch the videos *before* coming to class. Class time is spent elaborating concepts and working problems. While cynics would say that the students won't watch the videos beforehand, they will be totally lost during class if they don't. Furthermore, we think the students will watch videos (that include simple quizzes to test understanding) on their own schedule, making them feel more empowered.

4. TEACHING "DIGITAL SIGNAL PROCESSING" ONLINE

The course *Digital Signal Processing* will be offered on the Coursera platform in the spring of 2013, and has so far attracted 20,000 enrolled students. It is based on the "classroom version" offered by the Communication Systems department of EPFL to third year undergraduates. In the context of the standard curriculum at EPFL, the Digital Signal Processing course represents one of the first exposures to a higher-level engineering class for most students, who have spent the previous two years primarily on propaedeutic subjects. With respect to the studying habits of the attendees the course is clearly a game changer; the students are required to learn a set of tools which are mathematical and abstract in nature but which need to be applied to very tangible real-world problems. This is a drastic departure from the clean-cut and rather artificial world of, say, calculus or linear algebra textbook exercises. As a consequence, we learned early on that a key aspect of the teaching process is to persistently underscore the link between the mathematical abstractions and the applied side of the discipline and to provide as many nuggets of practical intuition as possible. In the classroom the main tools to achieve those goals are the ability to modulate the level of detail with respect to the immediate feedback from the audience and the possibility of impromptu and often informal digressions based on comments, incidents, and even current events. These staples of the standard teaching performance unfortunately translate poorly (or not at all) to the world of online teaching; this, we feel, is the main challenge that a MOOC syllabus has to overcome.

One potential solution is obviously to simplify the material and concentrate mostly on the applicative side. Our first operational decision, however, has been to preserve the academic level (and, consequently, the demand placed on the students) with respect to the class offered at EPFL. Obviously the fruition patterns for MOOC classes are radically different from the standard setup of university-level coursework. Firstly, the audience is much more diverse both in terms of expectation and in terms of prior experience; it would be futile to try and hit a middle ground that perfectly balances two contrasting requirements: providing extensive background while avoiding being boring. Luckily the medium comes to our rescue: we will simply rely on hyperlinks for the former and on the fast forward button for the latter.

In order to contextualize the sometimes arduous path to the goal, we instead deployed a two-pronged approach. The first we could call "putting the cart before the horse:" prefacing theoretical section by examples which, although not yet analyzable in formal detail, create a sense of expectation and necessity for what's to follow. Secondly, we tried to deploy "leitmotifs" in the form of practical examples that reappear in several places to highlight different theoretical aspects. These recurring themes are designed to create a sense of familiarity as new concepts are introduced.

From the organizational point of view, the course is divided in a series of "modules," with each module covering a specific topic such as Fourier Analysis, Filters, and so on; think of a module as a chapter in a book which, in our case, is the freely available companion textbook [2]. Each module is internally divided into a variable number of units and each unit adheres to the current MOOC practice of providing short, self-contained lectures with occasional interactive quizzes. Typically, one day of lecture contains three units and is followed by a homework. One of the nicest things about digital signal processing is that nowadays any personal computer is a fully functional DSP "lab;" students can easily implement and play with all signal processing tools via well-known commercial packages like Matlab or its free and open-source alternatives such as Octave, Freemat or Scilab [3–5]. In order to "close the loop" on the theory, so to speak, we provide one or more programming labs per module, with worked-out examples and exercises.

In terms of production we decided to embrace the acousmatic paradigm[3] and produce visually simple lectures in which a slideshow

[2] The platform allows images and supports LaTeX equations.

[3] "Acousmatic sound is sound one hears without seeing an originating cause [...] a term used to refer to probationary pupils of the philosopher Pythagoras who, so that they might better concentrate on his teachings, were required to sit in absolute silence while listening to their teacher deliver his

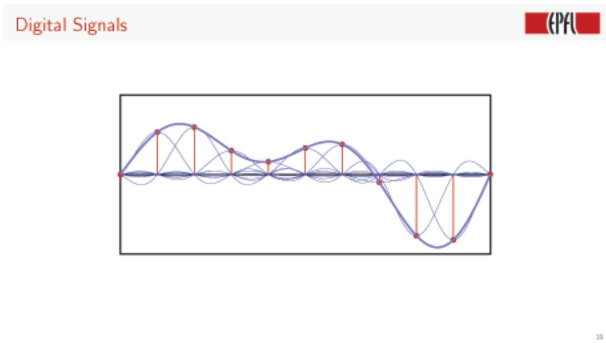

Fig. 3. A sample slide of the Digital Signal Processing class.

unfolds under a voiceover; annotations drawn by a "live" hand underscore the key passages (Fig. 2). This choice also allows us to record and edit the audio in a more efficient way prior to recording the annotations. To preserve the "human component" of the class, however, brief introductions and closing remarks are added to each module in the form of short videos in which the teacher appears "in person".

A remarkable advantage of the MOOC medium is the possibility of organically including animations, videos and sounds into a lecture, thereby easing the job of building bridges between theory and practice. To offset the sometimes abstract and dry nature of some mathematical concepts, great effort has been placed into the design of slides with as many illustrations as possible (Fig. 3). From a stylistical point of view, we strived for a simple and coherent graphical design and, to this end, we designed a LaTeX library, based on PsTricks, to efficiently draw 1D and 2D digital signal processing graphs in PostScript; the package is freely available at `http://www.sp4comm.org`.

A companion textbook for the course is available both in print and as a free PDF download [2]. The second edition of the book is currently being prepared for publication in electronic form only; many of the examples and of the animations developed for the online course will be included in the eBook in the form of interactive widgets. We feel that the synergy between the interactive eBook and a full set of video lectures will be extremely fruitful pedagogically.

5. EXPERIENCES AND LESSONS LEARNED

Teaching online has so far been an exciting and revealing experience. The ability to affordably reach thousands of students—often in regions of the world where high-quality education is simply unavailable—is, in our opinion, the greatest benefit of MOOCs. Students from almost all countries have enrolled in our courses, demonstrating the large demand for an advanced but affordable education. Creating a MOOC will also improve the quality of our on-campus teaching, rather than replace it. We are planning to experiment with the "flipped classroom" model, whereby our students watch videos at home at their own pace, and then come to class to discuss their difficulties, as well as applications and intuition behind the material.

In our experience, the main bottleneck for teaching online is producing the videos. First, it is not sufficient to recycle the materials from one's traditional on-campus course. Significant effort is required for: i) organizing and segmenting the material into 10–15 minute fragments; ii) delivering the material concisely, using the most meaningful examples and illustrations; iii) predicting beforehand the questions that the students will pose, and pre-addressing them in the video. Second, recording the videos in a clear, error-free, and professional fashion is extremely time-consuming. In our production workflow, every hour of produced video required roughly eight hours of preparation, recording, and editing. Since the videos cannot be easily modified after recorded, we often needed to re-record an entire segment just to fix some minor inaccuracies. We believe that lecture production can be made considerably more efficient by developing better tools that are specifically tailored to creating MOOCs.

Assessment is another strong limitation of MOOCs. While content delivery can be easily scaled to accommodate thousands of students, measuring the skills acquired by these students cannot be performed at scale yet. Accurately grading questions beyond the traditional multiple-choice or short answer format—such as long answers or mathematical proofs—is still not possible. We believe that peer grading may provide a viable solution for this in the near future.

Another problem related to assessment is that of certification and cheating. For example, it is fairly trivial for a student to create multiple online accounts, and repeat the exercises until they get perfect score. Or simply ask for external assistance while solving an online test. While cheating is not a major concern yet, surely, if in the future the course completion certificates become valuable enough, we can expect some students to cheat in every available way. The current solution (e.g., adopted by Coursera and edX) involves sending students to a physical test center, where appropriate measures are taken to verify identities and minimize cheating.

It is interesting to imagine how the teaching and certification will evolve in the coming decades. As more people learn from increasingly diverse sources—such as Coursera, edX, iBooks, and other yet-to-exist platforms—it is not unreasonable to imagine that certification/credentialing might separate from teaching, and perhaps be provided by independent parties. This contrasts with the current situation where traditional universities jointly provide teaching and certification. This separation of functions could, in principle, democratize certification, in the sense that skilled students who received their training from unconventional sources (e.g., online courses and articles) would have the same opportunity of demonstrating mastery as a privileged student who attended a prestigious university.

We are eagerly anticipating the feedback that will pour in once our classes are offered next spring. We believe that, as for any new tool, challenges will be abundant and adjustments will be required aplenty; also, we are well aware that many questions remain open, especially with respect to grading and certification. Yet we remain extremely optimistic and, for the time being at least, we find the trenches quite an exciting place!

All the results presented in this article are reproducible. Codes for generating the figures are available at `http://rr.epfl.ch`.

References

[1] D.H. Johnson and J.D. Wise, Jr., "A different first course in electrical engineering," *Signal Processing Magazine*, vol. 16, pp. 34–37, 1999.

[2] P. Prandoni and M. Vetterli, *Signal Processing for Communications*, EPFL/CRC Press, 2008, http://www.sp4comm.org/.

[3] "Octave," `http://www.gnu.org/software/octave/`.

lecture from behind a veil or screen." (source: `http://en.wikipedia.org`)

[4] "Freemat," http://freemat.sourceforge.net/.

[5] "Scilab," http://www.scilab.org/.