Assignments that matter

Once is an accident, twice is coincidence. Three times is a pattern. Over the past few months, I have bumped into repetitive evidence that the assignments I gave to my students were less than what they could be. Don't get me wrong: it's not that my exercises are flawed, unchallenging or uninspiring. I spend time preparing them, I work hard to balance between rewarding my students' comprehension and stretching their skills. I also like to make the instructions witty and fun to read. No, my assignments are ok, really. It is just that I started to realize that they – simply – miss the point.

Take this entry-level computer literacy class I have been teaching for the last 3 semesters at my institution. It's a required core class, which students have to take regardless of their technical expertise - needless to say, the majority of them are proficient with office and internet tools well before they enroll. However, I enjoy the class as a way to give students the liberal arts culture that's necessary to understanding where computers have been and where they're headed. We learn about some fairly intricate details of computer architecture, study the history of major players in the industry, failed attempts to introduce new technologies and so on. On this line, I spend one module talking about computer fonts. My usual assignment for this part of the class is one of my favorites: I assign to the class a few different well-known typefaces (Arial, Garamond, etc.) and ask students to research their history, how they were designed, who owns them, how they ended up in their favorite operating system. We then debrief in class, and inevitably, we are amazed and amused at the diversity of facts and trivia one can gather (did you know for instance that Arial was reportedly invented to avoid paying licensing fees for Helvetica?), as well as the depth of the debate involved when comparing the readability or mood of one font against another. This assignment, if not a difficult one, has always been a sure win.

Then it happened. One morning, in the electronic edition of my University's daily newsletter, I found a full-page feature on a professor in our School of Dentistry who, as a hobby, had published a study about the history of the little-known typeface used… in my University logo. Highly emblematic, charged with pride for alumni and simply irresistible for the rest of us encountering it everyday on our stationary or business cards, the story was relayed in the alumni magazine and strung a deep chord all around. How could I miss that one?!- I asked myself. Of course, I'll use that example in my class next time – again, a sure win. But just *how* could I not have the idea of assigning *that one font* to my students myself, in the first place?

I closed the email, sighed, printed the story and posted it well in sight on the board above my desk, determined to learn the lesson.

Then it happened again. This time, I was teaching a class about music signal processing – my research specialty. One of my assignments for the class is

meant to introduce students to the Fourier harmonic decomposition of musical sounds. I give them the recording of e.g. a guitar chord, and ask them to compute a Fourier spectrogram, identify frequency partials, associate them in harmonic series and deduct the actual notes composing the original chord – i.e. use computer analysis to produce a complete musical transcription. After some experimentation, I had come upon what I believed was the perfect example for the exercise, a well-known, pleasing, iconic sound for which a perfectly unambiguous solution existed: the opening chord of the Beatles song *Yesterday*, a simple F major.

The same week of January 2009, the Wall Street Journal reported about a Prof. Jason Brown, a Professor of Mathematics in Dalhousie University applying Fourier analysis to solve "one of the greatest mystery in the history of Rock'n Roll": the opening chord of… the Beatles song A Hard Day's Night. *That* chord had intrigued musicologists for years, and even George Harrison wouldn't exactly remember what notes he'd played. The story was published in a refereed mathematical journal, and relayed in Wired, Guitar Player Magazine and on National Public Radio. A Google query with "Beatles chord mystery" yields 320,000 results as I write these lines.

I closed, sighed, printed, posted on the board.

And I began to think.

Both examples were clearly better than mine. They had generated publishable work –achieving that crucial "leveraging of teaching towards research" that everybody says we faculty should pursue. They provide both great learning outcomes and certainly get the students excited. More importantly perhaps, they embody what a university education should be about: changing the world (even in a small way, be it only for Beatles fans). These are assignments that matter.

What do these share, that my usual exercises miss out?

- First, they operate on very *carefully chosen data*. Contrary to cutting edge research maybe, the innovation here rarely resides on the choice of method. Both me and Dalhousie's Prof. Brown were spot on for the Fourier thing – we both knew it'd work, and we wanted the students to learn about it. The big difference is choosing just the right song to apply it on. Too often, we spend great time polishing a careful illustration of a technique, and then generate dummy school-book data in the last minute – well, I say we stop doing that.
- Second, these are assignments for which the answer is not *known in advance*. The key is not to generate results that are useful (as would be e.g. the exercise of digitalizing some books), but to be the *first* to do so. A class being the first one to tackle and answer a question – however small - generates something psychologically distinct from a mere sense of "contributing". These are Indiana Jones kind of

assignments. Of course, in most cases, instructors could work out the answer first, on their own, but I believe this misses much of the pedagogical point (and the fun).

I'm now trying to develop a systematic eye for spotting such class projects. However I'm finding that this doesn't always fit well with what's expected of a university faculty.

First, it takes time and requires to be always on the look-out for data and opportunities. For instance, to teach number crunching in Excel, I used to assign students with a standard toy dataset of human resource records (fake employees of a fake company). I have since then started using real music sale charts published weekly by the Billboard company, and made available by a public API on the web. Planning for this one-day class exercise took weeks – I even had to learn a new programming language. Of all the advice I read about faculty time management, I have yet to find any that suggests this is good practice! Forget course planning, too. If Michael Jackson happens to die one week in advance, you cannot but reorganize the class schedule overnight to seize the opportunity to watch his record sales soar in real-time, from this unexpected front-row seat of yours. All this preparation and vigilance comes at the detriment of longer-term work, such as research.

Second, designing such assignments is a strange thought process, quite the opposite of research thinking. I was trained to look for the best method to solve a given problem (e.g. Fourier to analyze a chord), but now I need to find the best problem to use a given method (e.g. that opening chord for a Fourier assignment). Fudging with an assignment until it becomes "wow!", "way cool", until it "takes my breath away with its audacity"[1] (to paraphrase management gurus like Tom Peters) requires more project management skills than traditional scholarship. Indeed, part of that redefined job is a sales game to get the students excited, to seek collaborations when needed, feel opportunities or even markets ("the Beatles enthusiasts out there would love this"). More the work of a MBA than a Computer Science PhD.

Third, and this is worrying, I'm not sure how such efforts contribute for faculty evaluation. My institution is currently revising its AAUP-inspired faculty handbook, generating lively discussion about e.g. what characterizes effective teaching ("the ability to vitalize learning, to arouse curiosity, etc."). Sure, such assignments qualify, but exactly how much more than the "usual" ones? The clause that a class assignment should generate e.g. new knowledge is nowhere to be found. In fact, embarking students in an exercise which is necessarily less controlled than a known-answer test could even be considered irresponsible or unethical. Is it good research, then? Well, publishing a story about a computer font in the alumni magazine probably

[1] If "breath-taking" seems too much in this context, see how Harvard Law School Prof. Charles Nesson tasked his undergraduate class to take on the Defense in a Federal Court case (http://joelfightsback.com/, retrieved July 2009)

won’t get me anywhere closer to tenure, and most of such projects would be considered too limited in scope and/or incoherent as a whole. After all, they have to fit whatever class I have to teach, rather than a long-term research strategy. Service to community, then? A Beatles’ chord? Seriously...

Increasingly, I realize that seeing assignments in such a light is redefining my work, my scholarship and my relationship to my students and colleagues. While this is not always comfortable, my intuition says it’s both good for me, my students and my institution, so I’m going to keep going, and see what happens.

Jean-Julien Aucouturier is an Assistant Professor of Computer Science, on Temple University’s Japanese Campus in Tokyo, Japan. His website is http://www.jj-aucouturier.info