

Engineering Design I: Methods and Skills

Topic Readings

Steven H. Collins

August 27, 2015

Chapter 1

Learning Mechanical Design

When I was a Junior Mechanical Engineering student at Cornell University in the fall of 2000, I took a course called “Machine Design & Analysis”. It was a typical lecture course, with homework assigned from a textbook [Norton, 1998] followed by exams with similar problems. It was enjoyable enough, and I received an ‘A’. Later, as a graduate student, I chose “Design” as one of the subjects for my qualifying exams to pass into candidacy as a PhD student in Mechanical Engineering at the University of Michigan. The exam was based on questions from another textbook with similar format [Shigley; Budynas and Nisbett, 2006]. I passed. Life moved on, and I rarely considered these courses and exams afterwards.

In the fall of 2010, I became an Assistant Professor of Mechanical Engineering at Carnegie Mellon University. The department was looking for someone to teach a Junior-level course called “Engineering Design I: Skills and Methods”, a newer version of a prior course called simply “Analysis”. In the intervening decade I had become something of an expert in the practice of mechanical design through my research on humanoid robots and robotic prostheses. The topic of design was one close to my heart, and I was excited about the opportunity to teach it; I wanted to help new generations of students to become expert designers, as well as to share my knowledge of and passion for the subject.

I quickly realized, however, that I didn’t have a good model for teaching the topic. Traditional coursework had contributed little to my becoming an expert in design. Looking back over my undergraduate and graduate courses, there were few ideas or skills that had been acquired through listening to lectures, completing homework or taking exams. Most of my knowledge had come either from extracurricular activities, mostly research projects, or loosely structured, project-based courses that didn’t introduce core concepts (of mechanical design). I concluded that, at least for people like myself and at least on the topic of design, learning by doing was more effective than learning in the abstract through lectures and exams.

As it turns out, this experience is common to many people, and, more importantly, borne out by decades of research on pedagogy. Research has shown that courses based around projects, interaction, conceptual understanding and peer instruction are more effective than traditional courses based around lectures, memorization and exams. The differences are stark. For example, in introductory Physics courses, interactive methods more than doubled the amount of learning in a single lecture [Deslauriers et al., 2011], while peer instruction tripled the improvement in understanding over a semester [Couch and Mazur, 2001]. Double or triple the learning! These were techniques I was eager to try, and I have incorporated many of them into that Junior-level course on mechanical engineering design. This series of Topic Readings, for example, is used to transfer new knowledge without using precious classroom time, for reasons we will discuss below.

This first reading is about the structure of the course itself. Many students take some time to become accustomed to this style of course. Unfortunately, despite decades of well-regarded research, most university courses are still taught in a traditional lecture and exam style. Perhaps this is because such a format is comfortable and familiar, or because, despite resulting in less learning, it can be easier for both students and instructors. In any case, the change in style that students experience in this course can be jarring. This reading is intended to address many of the concerns that students commonly voice early in the semester by explaining the reasons that the course is laid out in this way.

1.1 Why are we doing things this way?

There are many elements in this course that differ from our fifteen years of experience in taking classes. Why did we do things this way? The short answer, and the organizing principle of the course, is: to better facilitate deep, real and true learning of mechanical engineering design. It may not be immediately clear how each choice relates back to this goal, however, and the following short essays are intended to help draw those connections.

1.1.1 Lectures

Why all the exercises in lecture?

Lecture is the most frequent topic I hear about in student feedback, both positive and negative. The above question is also sometimes asked as “why don’t we have a normal lecture?”, “why don’t we learn anything in lecture?”, “why don’t you

cover the topic readings more in lecture?”, or “why is there so much filler in lecture?”. The short answer is: although they require more effort from both student and instructor, in-class interactions and peer instruction are more effective at facilitating learning than the traditional lecture style of teaching.

Learning is hard work.

Wouldn't it be fantastic if learning were a simple matter of opening one's mind and receiving skill? My favorite version of this idea is imagined in the 1999 film *The Matrix*, in which the lead character has his brain plugged into a computer for a few moments, then announces “I know Kung Fu”. Unfortunately, many decades of research in neuroscience indicate that this is not how the brain works. Instead, neurons need to fire to rewire. That is, you have to think about something repeatedly over a long period of time in order for its imprint to be left on your brain. Thinking is hard work, best motivated by a challenging task that requires repeated application of the technique in question. Alas, we humans are not in the world of *The Matrix*, but rather that of *The Karate Kid* (think: wax on, wax off). We must practice to become skilled, and in-class exercises are one form of this practice.

How did you become an expert?

Consider for a moment something that you are skilled at compared to others. This might be something related to Mechanical Engineering, such as building racecars or designing in SolidWorks. It might be something from another domain of your life, such as playing an instrument, hitting a pitched baseball, dancing, or even playing a video game. Now consider the primary process by which you *became* an expert at that skill. Based on informal surveys of Juniors at Carnegie Mellon, it is likely that the answer is essentially *practice*, usually lots of practice, and often as part of a project, team or internship. It is unlikely that listening to lectures or reading a book was the primary activity involved in gaining this skill. (This exercise is borrowed from a lecture by Eric Mazur.) We must practice to become expert, especially in applied creative activities like design, and in-class exercises provide such practice.

Facilitating learning

The word ‘teaching’ seems to imply a process where most of the activity happens on the instructor side of the interaction. Since learning and expertise actually come from activities in the brains of students, the role of the instructor is more that of *facilitator*, *i.e.* providing the appropriate environment and motivation for learning to occur. Sometimes students provide feedback like “it seems like we’re expected to learn everything on our own”. Exactly so! There is no other way. In-class exercises are designed to provide such a rich learning environment.

Interaction versus information transfer

Class time is valuable, and must be used as effectively as possible. Exercises, discussion, peer instruction, and group critique require all of us to be in the same place at the same time. Information transfer, *i.e.* gathering the facts required for such activities, does not. In fact, information transfer has less and less to do with even being at school; with a quick purchase from Amazon or a few keystrokes at Google we can have almost any information we want (if we know what to look for). The traditional lecture format was developed in a different time. Producing and distributing a book in the 17th century was extremely difficult and expensive (think: horses). Even as late as a few decades ago, pre-internet, it was much more burdensome to find, say, a canonical equation or constant value without accessing a good library. Hence, the string of ideas stated by the lecturer were valuable information, worth storing in one's notes or even memorizing (later from those notes). That is no longer the case. In this course, we have moved as much information transfer outside the classroom as possible (see *Topic Readings*). This frees up class time for interactive activities, which are twice as effective at facilitating learning [Deslauriers et al., 2011].

Peer instruction

Professors have both strengths and weaknesses as teachers. We are knowledgeable in the subject area and have much experience in teaching it. However, there is usually only one of us to a large number of students, limiting one-on-one interactions. Perhaps less obvious is the fact that it has been a long time since professors learned the material in question, and we may have a hard time understanding the misunderstanding of a student. You and your peers have complementary strengths and weaknesses as teachers. In particular, there are better ratios and, having recently learned the material and coming from the same cohort culture, you may have a better sense of how to explain new ideas to your peers. As an added bonus, teaching someone something you've recently learned helps consolidate the understanding. Perhaps for these reasons, peer instruction during lecture has been shown to result in a three-fold improvement in conceptual understanding over the course of a semester in some courses [Couch and Mazur, 2001].

Group discussion and critique

Projecting randomly-chosen student work overhead for discussion and critique provides an additional enriching element to facilitate learning. It provides the professor with real-time feedback as to concepts at which students are excelling or struggling, allowing the ensuing lecture elements, examples, and exercises to be tailored appropriately. Group critique also provides similar benefits to peer instruction, but with a larger range of ideas brought forth. Such discussions and critiques are common in design as taught in a fine arts context.

Coverage versus understanding

As an instructor, it is tempting to try to fit more and more ideas into lecture on each re-teaching of a course. After all, we want students to learn as much as possible. Maybe if we just say each idea out loud once in front of everyone, the idea will permanently lodge in students brains? Then we could move quickly to the next idea and thereby say more different things out loud in a class period? This is the illusion of ‘coverage’, *i.e.*, that the number of things the instructor knows and mentions during class is a meaningful measure of student learning. What the instructor knows doesn’t matter, what the students know afterwards does. Passively listening to a long string of statements of fact does not allow a person to later recall all the stated facts, nor to later apply them. Pedagogical studies have demonstrated that students can only really learn a small number of new ideas in a lecture, typically fewer than about three per hour. Unsurprisingly, similar studies have shown that students’ skills in applying such ideas are much weaker following passive lectures than interactive ones [Deslauriers et al., 2011].

Imagine a dance course in which the instructor used all class time for their own demonstrations, then asked students to perform a routine as a final test. Because the instructor is skilled, they could perform many more maneuvers and routines during a class. But that would not be as effective as demonstrating fewer maneuvers and then having students practice and get feedback. Although in-class exercises use time that could be spent on longer lectures with more facts listed, in-class exercises are more beneficial because the fewer (most important) ideas that *are* discussed and practiced are actually retained.

Not because it is easy, but because it is hard.

Allow me to dispel the idea, implicit in some feedback I’ve received about this lecture style over the years, that exercises make a lecture easier to deliver. It is relatively easy for a knowledgeable professor to sit down with a textbook, transcribe a section of a chapter into slides or notes, and speak on that content. After all, we can talk on and on about almost anything! It is also relatively more predictable and less stressful to give a traditional lecture than to hold exercises. The problem is that it is simply not very effective at facilitating learning. The following quote attributed to Edwin Slosson (by way of Mark Twain) makes the point nicely: “College is a place where a professor’s lecture notes go straight to the students’ lecture notes, without passing through the brains of either”. To be even more cynical, sometimes it seems that there can be an implicit agreement that students will ignore the fact that the professor isn’t working too hard to teach if the professor ignores the fact that students aren’t working too hard to learn. But this is not what we will do; we will work hard to truly learn (and teach) mechanical design. One

of the key elements of this process will be practicing, discussing, teaching one another, and critiquing one another's work during class periods.

Why call on people randomly in lecture?

The short answer is that this facilitates participation from everyone. Lectures based on interaction will only be effective for everyone in the room if everyone interacts. When the professor asks a question to the room, one from a small set of students will nearly always answer, while a different set of students will nearly never answer. If we based our discussions only on comments made by those bolder students, the resulting lecture would be tailored to their needs at the cost of the needs of other students. If the boldest students are also working at a faster pace than the average student, this can hide the need for more practice in a particular area. On the other hand, the quieter students have a great deal to contribute, and asking for their input facilitates that interaction to the enrichment of the whole class. Calling on people randomly (based on a randomly ordered list of names) mitigates these and other biases that might otherwise emerge in participation and time allocation. Some students mention that this makes them worry they might be called on, which is unfortunate and of course not the intent. As a mechanical engineer thoughtfully engaged in an in-class exercise, you can always feel confident that your input will be valued and helpful to your peers, even if, perhaps especially if, it illustrates a misconception. On the bright side, you also don't have to worry about whether or not you should raise your hand, which was always a source of anxiety for me as a student.

Why do we cover things just before assignments are due?

The short answer is that that's when you're applying the ideas. Computers and hard drives can gather information and store it for as long as we want without losing anything along the way. Neuroscience tells us that this is not how our brains work. Our brains are constantly bombarded with huge amounts of data and therefore must be selective in what is retained. One of the primary mechanisms in this selection process seems to be 'use it or lose it'. Information that is immediately applied, *e.g.* in an in-class exercise, homework problem or project, is much more likely to be retained. In my experience, students, like most people, tend to work on their assignments just before they are due. We therefore try to organize the course to place topic readings just before corresponding in-class exercises, both just before homework due dates, and all of these things during a time in the project when such concepts would be most useful. This technique, known as Just in Time Teaching [Novak et al., 1999], has been shown to improve retention of concepts and application of concepts to new problems.

1.1.2 Topic Readings

Do I really have to read the topic readings?

The short answer is: yes! In courses taught in the traditional lecture style, keeping up with readings is usually not as important. Lectures often significantly or completely duplicate the textbook, and one can always go back to the book to cram just before an exam. In a course based around in-class exercises and discussions, reading things before class is crucial. Topic readings are the primary means by which new information is communicated, which cannot be made up for by attending lectures that synthesize, rather than repeat, the content. Topic readings are like the injection of fuel and air into an engine cylinder, exercises in class provide a spark, and bang! Learning. If one isn't primed by the topic readings prior to class, the exercises have much less impact and a learning opportunity has been lost. Reading the material later will not make up for the deficit. It would be like participating in a book club discussion and then reading the book.

Why are the topic readings so short/long?

Providing just the right amount of information, and nothing more, is a challenge in technical writing. Given the availability of additional information, e.g. in optional texts [Budynas and Nisbett, 2006, Dieter and Schmidt, 2009, Steif, 2012], we have tried to provide only the most essential information in these readings. Fortunately, we receive about equal requests for longer readings as for shorter ones, suggesting that they are around the right length.

1.1.3 Projects

Why do we have projects?

Just kidding. Everyone loves the projects and they require little additional explanation. Suffice it to say that they provide the *raison d'être* for the course; why learn how to design mechanical things except to actually design them? Projects are intended to require application of the most important knowledge targeted for learning in each week of the course, and to have a clear connection to real-world scenarios in which a mechanical engineer might find one's self.

Why won't you tell me what to do in my project?

The second most common subject of feedback in this course relates to "vague" answers to questions about projects. Students struggling with some aspect of their design will sometimes ask a question similar to "what should I do for this part",

and receive an answer similar to “when confronted with situations like this, try using x approach”, where x is a broadly applicable design technique, or be asked to talk through their high-level approach to the problem. Sometimes, a student will seem to prefer to be directed that “the best solution here is y ”, where y is some design element. In those cases, a student might then become frustrated and wonder “why won’t you just tell me what to do?” The short answer is: this course is about *how* to solve design problems in general, not *what* equation or design element would be good in one scenario.

Exams versus design projects

In a traditional lecture-exam course in a fundamental, rather than applied, area of mechanical engineering, we are often told which equations to use, under which (often contrived) circumstances, asked to (temporarily) memorize the equations and circumstances, and then to recall them for homework problems and exams. In such classes, students often ask professors questions of the form ‘how do I do x ?’, and receive an answer of the form ‘you should always take equation y , then use equation z ’, where ‘ x ’ is one of the canned problems and ‘ y ’ and ‘ z ’ are each one of the formulas expected to be memorized.

Design is different. There is no set of formulas that, once memorized, answer all the design problems you might encounter. Rather, design is about a process in which (mostly existing) knowledge is applied to real-world, creative problems. The set of equations that might be useful in such a process are really for *analysis*, or evaluating a candidate solution, rather than *design*, or creating a new one. Each design problem requires use of different analyses, drawn from a vast set of possible equations and tools. An important part of the process is selecting amongst the available tools. The best design elements revealed through iterative application of such analysis are also unique to each design problem and are drawn from an even greater set of possibilities. We facilitate students’ learning how to navigate the design process by providing guidance as to what high-level processes to use but not what precise equations or answers to arrive at.

Design as mystery

Arriving at a good solution to a mechanical design problem is analogous to solving a riddle or mystery, making a legal case, performing scientific research, or any other application of creative thinking, logical reasoning and the scientific method. We begin in a state of ignorance and possibly confusion, having only guesses as to what answers will be revealed. We consider all the possibilities we can imagine, and pursue the consequences of each analytically, iteratively refining our guesses. We seek out new information or perform empirical tests as needed. After much thinking and testing, we develop an understanding of the key parameters of the

design and the best ways we can conceive of addressing them. While answering some such mysteries as examples can be helpful, say through in-class exercises or in *post-hoc* project analysis, simply solving the mystery of a project for a student before they've had adequate opportunity to unravel it themselves would rob them of the most useful form of practice possible in design. Instead, we try to maintain something approaching real-world conditions; when working as a design engineer, there is no 'right' answer, nor oracle from which to seek it. We are on our own. Fortunately, we have all the tools we need to be successful.

Sometimes giving the desired answer isn't helpful

Answering questions like "what should this aspect of my design look like" would therefore defeat the purpose of the course. This would be analogous to answering the question "what is the numerical answer to this homework problem". Knowing that the best solution to one specific design problem is such and such is equally irrelevant as knowing that the answer to one specific exam problem is forty-two. We would never ask that question about the homework (until reviewing the assignment after its completion), but would instead ask the question of how to arrive at the right answer. In design, the analogous question to ask is what design process should I use to arrive at a good design solution. The following old adage is relevant here: give a person a fish and they will eat for a day, teach a person to fish and they will eat for a lifetime. Each design problem you will encounter in the real world is unique and complex, with wide variations in constraints and outcomes of interest. The high-level processes and techniques we use to solve them, however, are similar across a wide domain in mechanical design. Even though a student might ask a direct question, it doesn't mean that they should, for best learning, be given a direct answer of the desired form.

Specific projects for broader goals

Course projects, by necessity, address one specific design scenario, but they are a vehicle for addressing the much more important goal of learning how to perform mechanical design in general. It might sometimes feel as though the purpose of the course is to learn how to make a good, *e.g.*, astronaut's hat rack. After all, we spend a lot of time working on that particular project and discussing outcomes. If that were the purpose, it would make sense to simply show a good design on the first day of lecture and be done. But it isn't our purpose. As we've discussed already, the best way to learn a skill is by applying it. The project scenarios allow us to apply our skills, but we should be careful not to make the mistake of thinking that these particular answers are the point of the class.

Concept design phase

In some projects, we will enforce an initial concept design phase by selectively

withholding quantitative data that is not necessary to make solid initial progress on the project but might lead to distraction. This can also be frustrating to some students, perhaps because of the sense of solidity that comes from having hard numbers to write down. In such cases, I urge you to trust the process. Qualitative models and symbolic derivations can take you much farther in the design process than you might initially realize, and much faster than when bogged down in particular numbers. Keeping your process in symbolic form also frequently reveals surprising relationships that would otherwise be hidden in the decimals.

Not for my amusement

It is unpleasant for a professor to have to give a student an answer other than the type sought. We are empathetic with the frustration such a student can feel, and, of course, dislike the possibility for social conflict created. It would be far easier and less stressful to simply solve the mystery immediately. Unfortunately, as we've elaborated upon above, this would not be in the long-term best interests of the student. Please do not take it personally if your question isn't answered in the form desired. Rather, try to realize that the professor is working harder, and enduring the same difficulty and frustrations you might, to guide you through the process that is likely to solve not only this design problem, but many other design problems you might later encounter.

Why is part of the project grade based on objective performance?

The short answer is that this avoids over-reliance on the subjective opinions of the professor. Subjective evaluation of a design is tricky. It is inevitably based on the intuition, experience and mental models and projections of the person making the evaluation. Given my own, unavoidable biases, I might like one person's approach and not realize a key flaw or dislike another person's approach despite its being perfectly well suited to the particular circumstance. That wouldn't be fair. For this reason, I try to limit my subjective evaluation of projects to the *process* used, rather than the *product* created. Of course, in the real world, the primary outcome of interest is the product. What to do?

We address the need for evaluation of the design product itself through a uniform, objective measure of performance obtained from empirical testing. The equation and constraints are laid out in advance so that everyone has the same objective. It is difficult to know what will constitute strong or weak performance in advance, of course, so some of the terms in this measure must necessarily be set *post hoc* based on the designs produced by each student. Some students report feeling anxiety about testing of their components or machines. This is unfortunate, and is not the intent of including such an element, but seems less troublesome than the

alternatives. If the same portion of the grade were based on the less-predictable subjective judgment of the instructor, for example, that might be even more nerve wracking. At least the objective measures of performance are subject to the immutable laws of Physics, which makes them a bit more predictable.

Why don't you force us to follow a set project schedule?

The short answer is that this allows development of project management skills. Good time management and project planning are more important in design than in more theoretical areas of Mechanical Engineering. This is because good results are achieved only after many iterations that span days and weeks, and there are fewer ways to 'cram' such activities. Learning how to keep one's self or team on an effective schedule, with minimal external enforcement, is therefore critical.

In conversations with students who have found the project execution time to be too short, we have often come to agree eventually that front-loading their design activities more would have resulted in a more effective process. Students who find the projects to last too long are often skilled practitioners of mechanical design, who often later agree that they might have learned more from the project if they had spent more time refining their design, which was acceptable, towards a more optimal, highly-competitive one. Fortunately, the suggestions are fairly equally divided between too short and too long, meaning we're probably near the best we can do for a large class.

1.1.4 Aftermath

Although many students are at first surprised by the format of this course, most eventually come to enjoy it. Over the years, I've been glad to hear from many students that this was their favorite course at Carnegie Mellon, and pleased to hear from others that it prepared them well for their internships and jobs following graduation. I love teaching design, and look forward to a challenging, fun, and transformational semester together.

1.2 Acknowledgments

Thanks to the students, graders, Teaching Fellows, Course Assistants and Teaching Assistants in Engineering Design I: Skills and Methods (24-370) at Carnegie Mellon University from 2010 through 2015 for sharing their thoughts and feedback, which are the basis for the essays in this chapter.

Bibliography

Richard Budynas and Keith Nisbett. *Shigley's Mechanical Engineering Design*. McGraw-Hill, 8th edition, 2006. ISBN 0-390-76487-6.

Catherine H. Couch and Eric Mazur. Peer instruction: Ten years of experience and results. *American Journal of Physics*, 69:970–977, 2001.

Louis Deslauriers, Ellen Schelew, and Carl Wieman. Improved learning in a large-enrollment physics class. *Science*, 332:862–864, 2011.

Geogre E. Dieter and Linda C. Schmidt. *Engineering Design*. McGraw-Hill, New Jersey, 4th edition, 2009. ISBN 0-1070283703-2.

Robert L. Norton. *Machine Design: An Integrated Approach*. Prentice-Hall, Massachussetts, 2nd edition, 1998. ISBN 0-13-017706-7.

Gregor M. Novak, Andrew Gavrín, and Christian Wolfgang. *Just-in-time teaching: Blending active learning with web technology*. Prentice-Hall, 1999. ISBN 0130850349.

Paul S. Steif. *Mechanics of Materials*. Pearson, New Jersay, 1st edition, 2012. ISBN 0-13-220334-0.