

1 Introduction

For centuries university students have been listening to lectures from their professors. Clearly that approach can lead to positive outcomes, as students graduate knowing more than when they arrived and they go on to have success in their future careers. But is lecturing the most effective way of helping our students acquire the knowledge and skills they need to be successful scientists? The answer seems to be no (Freeman et al., 2014). Cognitive science research and classroom studies show how learners must actively engage with the material to gain a deep understanding and acquire the skills to successfully apply their knowledge in varied situations (Brown et al., 2014).

Active learning is not a new idea. We have all used it, for example to learn how to play a musical instrument, shoot a basketball, knit a sweater, or to ride a bike or drive a car. We also employ active learning in field trip and lab activities in our courses, where students typically learn to identify fossils by actually looking at fossils. However, active learning exercises still require that the students have some foundational knowledge, including a common vocabulary and familiarity with underlying concepts.

One approach to increase the amount of active learning in the classroom, while recognizing the continuing need to deliver content, is the “flipped classroom” model (for a summary, see Milman, 2012). The flipped classroom gets its name because it inverts the traditional lecture approach (it’s also sometimes called an “inverted classroom”), delivering content outside of the classroom and allocating in-class time to active learning exercises. The out-of-class content delivery phase of the learning cycle can take several different forms, and there are similarly many choices for how to structure the in-class activities.

In this Element, I will discuss my experiences at the University of California Santa Cruz (UCSC) in converting three lecture-style courses to flipped classes. I will focus both on benefits and also potential concerns of the flipped class format, and I aim to provide some practical suggestions for overcoming those challenges based on my personal experiences. These courses (Invertebrate Paleobiology, Sedimentology and Stratigraphy, and Statistics and Data Analysis in the Geosciences) are all upper-division electives taken primarily by juniors and seniors, with enrollments in the range of 15–30 students. As a result, my experiences may be different than if the courses were required, had larger enrollment, or were taught at a different academic level. UCSC is a medium-sized and moderately selective public R1 research university. In fall 2016, 42 percent of undergraduates were first-generation college students and

39 percent were Pell Grant recipients (www.universityofcalifornia.edu/infocenter/fall-enrollment-glance). UCSC is also a Hispanic-Serving Institution and, in fall 2017, 31 percent of Earth and Planetary Science majors were Hispanic or Latino. Our program is large (for an Earth Science department): 198 majors in fall 2017 and on average 60 graduating students per year.

2 Implementing the Flipped Classroom

First, I'll provide some background to describe how I have implemented active learning and converted my courses to a flipped class structure. There is no single "correct" way to flip the classroom so, although I will discuss my choices, it is important to craft policies and procedures that align with your personal teaching style and with the resources and time available to you. Flipping the classroom can go as far as shifting all content delivery to outside of the classroom and spending all in-class time on activities, but that is simply one end of a continuum. You don't have to go from 0 to 100 percent at once. I transitioned the invertebrate paleobiology course to a flipped class with an incremental approach, by initially adding in-class activities to some class meetings, so in those select class meetings we would spend 75 percent of time on lecturing and 25 percent of time on hands-on work. Over time I started recording videos to be able to expand the short activities into hour-long exercises that occupied the entire class time, and started designing more activities for other topics. In contrast, I shifted my sedimentology/stratigraphy course from pure lecture to 100 percent flipped class all at once. For reasons of instructor workload (and sanity) I would not recommend that approach!

To introduce the content, I produced short videos by recording my voice over PowerPoint slides (following a pre-prepared script), using PowerPoint's built-in slideshow recording option. The videos are typically around 10 minutes long, and no more than 15 minutes. I saved the videos as MPEG movie files and then uploaded them to a YouTube channel for the class. These YouTube videos are available for anyone to watch, but it's possible to make the videos unlisted and only available with a link you provide to your students. My approach was fairly low-tech and didn't require video-editing software or any special skills on my part, so it should be a process that anyone can adopt. Having a pre-written script helps the video narration flow more smoothly, but the script could also be posted for students to view. I have not tried this, but there are potential benefits, such as providing a framework for student note-taking, aiding accessibility for international students or English-language learners, and demonstrating the correct spelling of technical terms (where YouTube's automated captioning often fails spectacularly). If you

have the desire, there are more sophisticated possibilities for interactive online instruction, but they may require special software or technical support from trained personnel. Although the videos have been surprisingly popular among students, videos or online content are not the only choices for pre-class assignments. To give some examples, you could assign readings from a textbook or the internet, have students read an article from the scientific literature, or ask them to participate in online discussions.

My in-class activities are primarily pen-and-paper worksheets with conceptual exercises, the kind of exercises that many of us assign as homework in a lecture-style course. Some involve hands-on work: For example, students flip coins to generate a bounded random walk, or explore adaptations to soft substrates using metal and clay shapes of different sizes in containers of water-saturated sand (this activity can get a bit messy . . .). Others have students analyze real or simulated data and models on computers during class. The full list of exercises can be found at my website: <https://people.ucsc.edu/~mclap/ham/eart101.html>. I have also found inspiration for activities at the SERC website (Science Education Resource Center; <https://serc.carleton.edu>), which is a great resource for lessons that can be used or adapted for in-class work. My primary suggestion when designing in-class exercises, such as the design of pre-class videos or readings, is to work within your comfort zone and choose the types of activities that match your teaching style.

The transition to a flipped classroom can be a stepwise process, so it is beneficial to experiment, reflect on what worked and did not work, and then adjust. This advice applies not only to the in-class activities but also to pre-class content delivery. That is not to say the shift is without challenges, but adopting a flipped classroom doesn't require any special training beyond the motivation to try. The challenges associated with the flipped class format can be overcome, and the benefits are well worth the effort.

3 Benefit: Students Learn More

The most fundamental benefit from switching to a flipped class structure that incorporates in-class active learning is that students learn more than they would in a traditional lecture (e.g., Moravec et al., 2010; Deslauriers et al., 2011; Mason et al., 2013; Freeman et al., 2014), although the improvements may derive largely from active learning rather than the flipped structure itself (Jensen et al., 2015). The improved performance from active learning may be especially pronounced for underprepared students (Ryan and Reid, 2016) and those who come from more disadvantaged backgrounds (Haak et al., 2011; Freeman et al., 2014), although not all studies support this (He et al., 2016).

My anecdotal experiences have been consistent with these research findings. Although the flipped classroom isn't a cure-all, it has been noticeable how students do in fact gain a deeper understanding of the material with active learning. I made simultaneous changes to exam rules and structure that prevent a more quantitative assessment, but my qualitative impression is that there has been particular improvement in students' ability to apply their knowledge to more authentic problem-solving situations. The flipped class format allows (more accurately, forces) students to grapple with actual problems, which in many cases can be open-ended and partially ambiguous, providing them with a more realistic experience. Rather than focusing on superficial memorization, students actually gain skills that scientists use. I've been pleasantly surprised by the difference that shifting to a flipped class has made to student confidence and expertise.

3.1 Associated Challenge: Students May Not Recognize That They Are Learning

Because the in-class activities often feature authentic problem-solving tasks, which forces students to confront challenging and unknown questions, students can feel like they are struggling and therefore that they aren't learning. Of course, this is the opposite of the truth! This misperception likely arises because many students don't realize that effortful practice on challenging material and retrieval of information by self-testing actually promote learning (Brown et al., 2014). Instead, people often fall into the illusion that learning occurs when something seems to come easily. A well-presented lecture can make the information seem straightforward, which seduces listeners into believing they have mastered the material (the fluency illusion). There is also sometimes concern among instructors that students' perceptions that they aren't learning as much may lead to more negative student evaluations of teaching. What strategies are there for conveying the benefits of active learning to students?

I'll start by noting that I haven't seen any drop in numerical scores on student evaluations of teaching aside from normal year-to-year fluctuations, and positive written comments about the flipped class format far outweigh negative ones. That said, I have received two instances (over seven flipped class offerings) of strongly negative written comments from students who vehemently disliked the format! I suspect there are also students who dislike the format but who do not complete the student evaluation. I think the concerns about significant negative consequences for instructor promotion/tenure cases are generally overestimated, but I recognize that student reactions may be different

depending upon the circumstances and that the criteria by which faculty teaching is reviewed will also differ among institutions. Although my experience has been positive on balance, I have attempted several strategies to further convince students of the benefits.

I began to include a paragraph in the course syllabus that briefly outlines the cognitive science justification for active learning, and I also try to relate instructional design choices back to those principles. If students mention feeling confused or remark that the exercises are difficult, I may emphasize how deliberate and deep thinking is actually a good sign of learning (of course, I also help them clarify their reasoning to move beyond the specific point of confusion). I also try to insert discussions about topics such as the importance of spaced practice, how repeated retrieval of information promotes long-term retention, and how knowledge is more easily acquired when building upon a foundation of prior information. I also emphasize how science, and education more generally, is primarily not about learning a set body of facts, but instead about acquiring the skills to interpret data, solve problems, and apply knowledge in new situations. My impression is that students are generally interested in learning about cognitive and metacognitive strategies, and that these informal discussions do help limit student resistance to active learning.

More recently, I've added questions about student goals and expectations to the questionnaire I distribute on the first day of class. The questionnaire was initially intended to learn about students' preferred name and pronouns and to gauge their relevant prior coursework, but I now also ask two additional questions. The first has the students think about a hobby, sport, or other skill and describe what they did to become proficient at that activity, with the aim of raising the topic of deliberate practice (and sometimes struggle) as an important way of learning. The second asks the students to rank three learning goals in order of personal importance: acquiring information, learning how to use information and knowledge in new situations, and developing lifelong learning skills. Most students rank the latter two goals more highly, which allows me to raise the benefits of active learning in developing skills to apply science as a process rather than as a body of facts. The goal of these prompts is to help prime students to view rigorous effort as a positive trait, so it's possible they lead to some improvement in student attitudes.

4 Benefit: In-Class Work Can Be Made Challenging

The flipped classroom isn't the only way for students to work on exercises; it's also possible to design high-quality homework assignments in a lecture-based

course. So why go to the trouble of flipping a class? Although homework assignments also give students opportunities to practice interpreting data or applying the concepts from lecture, I think it’s more feasible to implement more open-ended and possibly ambiguous exercises when instructors are on hand to provide guidance. Many of the activities I use in class, especially in the sed/strat course, are more challenging than I would assign as homework. Working in class with instructors present gives students the confidence to extend themselves and tackle these challenging problems.

4.1 Associated Challenge: Maintaining Student Engagement and Progress

Students only gain the full benefits of active learning and distributed practice if they complete the activities in the sequence and at the time at which they are assigned. However, students have many demands on their time and may not always be able to prioritize your class. How then to best ensure students participate in and are engaged during the in-class work?

I have experimented with several methods of grading to incentivize students to complete the in-class work promptly (Table 1). I initially thought of the classroom activities as in-class equivalents of traditional homework, so I

Table 1 Grading strategies for in-class work

Strategy	Benefits	Drawbacks
Grade based on correctness of answers, exercises turned in at end of class meeting	Incentivizes student effort and promotes timely completion	Highly stressful for students learning material for first time, distracting from learning
Grade based on completion, exercises turned in at any time	Allows students to concentrate on learning and provides flexibility with other demands	Some students let work pile up, negating benefit of spaced practice
Grade based on completion, full credit for exercises turned in by next class, partial credit for at any time later	Provides flexibility when students face other demands, but signals importance of spaced practice	My preferred option, although the incentive for timely completion may not be that strong

graded student work based on performance. Students ideally turned in the exercise at the end of the class meeting, although I sometimes allowed assignments to be turned in at the next class when I misjudged the time that would be required. This approach, in theory, should be very strong motivation for students to attend class and complete the exercises in a timely fashion. However, even though I was deliberately very lenient, assessing performance was still stressful for the students, so it ended up distracting from learning. In hindsight, mixing “teaching” and “testing” was a mistake.

I quickly shifted to grading in-class work for completion only. Assigning points only for completion of the exercise reduced anxiety and allowed students to concentrate on learning, but policies governing the completion of those exercises are still important. Until recently I encouraged students to turn in completed assignments at the end of class or the beginning of the next meeting, but I would give full credit for assignments turned in at any point. The flexibility of this policy was appealing, especially when students had circumstances preventing them from focusing on the class, and I also reasoned that it was better for a student to complete the exercise much later than never at all. However, in rare instances students would fall far behind and then complete a large number of exercises all at once, defeating the purpose of the distributed practice inherent in the course design. After this experimentation I have settled on a compromise approach that assigns full credit for completion of the exercise by the subsequent class and partial credit for completion of the exercise at any time. The difference in points between on-time completion and completion at any time is small enough that the effect on a student’s final grade will be minimal, but this system signals the importance of staying on top of the work.

Accommodating the pace at which students work through the exercises is an additional challenge, as some students will pick up the concepts more quickly while others may require more time or will be more thorough in their work. I have found this one of the most difficult logistical issues to address. One possible mitigation is to include additional questions that are not part of the assignment but that students can consider if they have extra time, as a simple form of differentiated instruction. I have built these additional questions into several in-class exercises (although not all of them), and students generally will take the time to work on the extra material. The interactive format also allows me to provide more guidance to students who need support and to let stronger students struggle a bit with the material before I help. I also partly avoid the problem altogether because the in-class exercises tend to run a little long, although that’s not ideal. Despite those mitigation strategies, at times some students finish a few minutes early – typically not more than 5 or 10 minutes

before the end of the 65-minute class period. If students finish the material early, I don't force them to stay and they're able to leave class. Although this may seem like anathema to professors, and may be problematic if student contact time is monitored by administrators, in principle teaching should focus on student learning rather than time spent.

4.2 Associated Challenge: Accommodating Greater Instructor Workload

Frequent in-class work is valuable for student learning, but the process of giving feedback can impose a significant burden on the instructor's time. I will begin by noting that my teaching load involves 3.25 hours of contact time per week over a 10-week quarter, so strategies that worked for me may not be feasible for someone who teaches more. What strategies are there for maintaining the quality of feedback to students but also reducing the burden on the instructor?

For the first few years of teaching flipped classes, I provided detailed and individualized written feedback on each student's assignments, which hopefully was valuable to them but was extremely time-consuming (1–3 hours of work per assignment) for me! Although feedback is useful in theory, it wasn't clear that students read the comments in any detail, so I have more recently begun to explore other options (Table 2). Delegating the assessment to a teaching assistant or grader could be feasible at some institutions, but I didn't want to add to the workload of our TAs, who already help out during class, teach the lab sections, and come on field trips. It may also be difficult to expect TAs or especially undergraduate graders to provide high-quality feedback, given the range of students' answers to open-ended questions. My latest strategy eliminates individualized feedback and substitutes a detailed key that I make available to the students after the exercises have been submitted. I collect student assignments to scan for widespread misconceptions that we can discuss as a group, but otherwise I don't make any annotations to the students' work. Although this has reduced the assessment workload dramatically, my initial feeling is that this approach is a little too impersonal and that it is also possibly not clear to students whether their answers meet or fall short of expectations. In future class offerings, I plan to adopt a hybrid approach, providing the answer key but making notations (not written comments) on each student's assignment to mark the answers that fail to meet expectations.

There is some initial time commitment to set up the keys, which require more thorough written explanations than would be necessary for my own grading

Table 2 Strategies for feedback on in-class work

Strategy	Benefits	Drawbacks
Instructor provides individualized, written feedback	Granular and detailed information for students	Extremely time-consuming
<i>Teaching assistant provides individualized, written feedback</i>	<i>Would reduce burden to instructor</i>	<i>Teaching assistants likely have high workload already</i>
Instructor makes answer key available	Detailed information provided to students with little burden on instructor	Impersonal, perhaps not clear to students if their answers meet expectations
Instructor makes answer key available and indicates when student answers fail to meet expectations	Detailed information provided and semi-personalized feedback	None? I think this is the best compromise and it's the approach I'll adopt.

Italics indicate that I have not tried this approach.

purposes, but there will be time-savings over the longer term. This would be especially true for larger classes. However, keys have one potential risk in that they can be passed on to students who take the course in a future offering. If work is entirely in class it is difficult for students to refer to a previous key, but if students are able to complete work outside of class there is a risk that they may instead copy the answers. There is no foolproof way of identifying this sort of academic dishonesty, but answers copied from a key are likely to be unusually similar to the key answers themselves, a good clue especially if the activities are more open-ended or realistic. Providing individualized and detailed written feedback is an aspirational standard, and may be achievable depending on instructor workload and/or availability of teaching assistant resources, but my experience suggests that the use of answer keys is a compromise that maintains most of the educational benefit.

5 Benefit: Students Work More

One lesser-publicized benefit of a flipped classroom, beyond the learning gains from active learning strategies, is that students often must spend more time working on class-related activities. This will not always be the

case (e.g., He et al., 2016), depending on the design of any particular class, but in my experience it's rare for a traditional class to assign homework assignments after each meeting. In contrast, the flipped classroom typically requires students to complete a reading or watch a video before every class. For example, I tell students that they should plan to spend 30–45 minutes watching the pre-class video (the video itself is 10–15 minutes), taking good notes, and ensuring that they understand the concepts. Although I don't believe many students actually spend 30–45 minutes, any amount of out-of-class effort is on top of the in-class contact time.

5.1 Associated Challenge: Ensuring Student Compliance with Pre-Class Work

Pre-class content acquisition is an important foundation for the in-class activities, but do students actually complete the assigned pre-class work? The potential that students may not be adequately prepared is a concern for instructors, so what strategies are best for ensuring or incentivizing compliance? Do these strategies motivate students to engage meaningfully with the pre-class work?

First, I'll note that my experience in upper-division elective courses is that students are generally conscientious, although there certainly are times when some students watch the video just as class is beginning or fail to watch the video altogether. The experience may be different for a class that is required or in a general education class designed for non-majors, as some students in those courses may have less motivation for the subject material. That said, I have experimented with different strategies for pre-class compliance (Table 3).

Initially I tried low-stakes online quizzes administered through and automatically graded by our campus learning management system. Most students completed the quizzes, but it's not clear how deeply they engaged with the material because the quizzes were designed to be quite easy. The quizzes also comprised only 10 percent of the overall grade, divided among 20 or more quizzes, so the incentive to complete any single quiz was not that strong. One additional downside was that students would occasionally claim technical difficulties – for example, that they submitted the quiz but the scores didn't appear in their record – and those claims were difficult to adjudicate. It may be less intrusive to integrate assessment questions with the content delivery in a more sophisticated online design, but my use of voice-over PowerPoint videos didn't permit that approach.

I wasn't particularly satisfied with the use of quizzes as a penalty for non-compliance, so I have also experimented with an honor-system approach