

Reading Fast for Fun* and Profit: *Massive Reductions in Stress, **Better Student Feedback**

Dr. Joshua Earle, University of Virginia

Joshua Earle is an Assistant Professor at the University of Virginia in the Department of Engineering and Society. His work focuses on the history of eugenics and its legacies in contemporary medical science and technology, including assistive technologies and artificial intelligence.

Reading Fast for Fun* and Profit**

*Massive Reduction in Stress **Better Student Feedback

Introduction: aka The Problem

Tell me if this sounds familiar. It was my first year of a full teaching load and I was staring down the gullet of ninety capstone Science, Technology, and Society (STS) research papers. History and English professors will certainly scoff at such a meager number of papers (and *only* 4,000 words each? pssh), but giving written feedback on those papers destroyed my life for several weeks. My colleagues commiserated, many feeling the same, but had little in the way of solutions except to have a stable of go-to feedback which you could copy and paste into Canvas to shorten the time spent.

I tried that the next year... with limited results. Along with finding little functional improvement in students' papers after giving written feedback (they would mostly just change exactly what I asked, and nothing else, even if I told them that there were more instances), I had to find another way. I refused to give my life away to grading for several weeks every semester, burning myself out and failing to actually shift students' skills in any positive way. Work-Life balance is real, and even academics deserve it.

So, I timed myself reading a full paper, and how long it took me to complete the written feedback. I noticed I could read the paper in about ten to fifteen minutes, but that gathering my thoughts into writing took me at least another thirty, and then I had to disconnect and wait until I could focus on the next piece. The whole process took me about an hour per paper. Which meant 90 hours of grading. That tended to take me about three weeks once I factored in the rest of my job responsibilities. It would bleed into my after-work hours, when I should be trying to relax and have a life outside of my job. And the stress of having that task like a Sword of Damocles over my head, meant that my brain wouldn't let me relax, wouldn't let me decompress, wouldn't let me pay attention to a book or TV show. And as someone without children or family responsibilities, many (if not most) of my colleagues would be facing even greater pressures to their time.

But it only took me ten to fifteen minutes to read a paper and understand where the major problems were. Instead of writing my thoughts down, I wondered if I could just say them aloud? How long would that take? What if, I thought, I just brought my students into my office, read their piece while they were there, and then told them what I saw? Could I cut the amount of time I spent giving feedback down at all? If so, by how much? What effect would it have on student learning and improvement?

In this paper, I will break down my class situation, how I implemented these changes, and what the results have been for me so far. I will then go into the broader research on pedagogy and the alternative systems that I am taking, and follow that with an acknowledgement of the limitations of my system. I hope that my experience can be insightful or useful to other educators, especially as class sizes rise and responsibilities mount in an age of austerity and external attacks.

My Situation

My main teaching obligation is a 2-semester capstone course series for engineering students, STS 4500 and 4600. In this series, students learn how to plan and execute an STS research project, write an original research paper of publishable quality, and produce the majority of their graduation portfolio. I teach three sections of twenty-eight to thirty students each semester (depending on enrollment and instructor availability), with the majority of the students staying with me through both courses.

I am not the only one teaching these courses. In fact, most of my department's teaching is focused on these courses and our 2nd year Engineering Ethics course. This spring we had twenty six sections in total, taught by fourteen instructors. Each instructor has a lot of leeway to teach in their own way, though many of the main learning objectives, and certainly the main deliverables each semester are the same across all sections.

So, I (and everyone else) have to give feedback on around ninety papers per semester, often multiple times. I set my semester up so that I have three drafts before a final version. The first two drafts are actually four, but don't tell my students that. My first draft is due very early in the semester (earlier than anyone else in the department), as I want them to be done with the difficult part of putting words to paper quickly, and move on to the more important and more worthwhile task of revising and polishing. The first draft is then peer reviewed during class and

over a weekend. The students have a week to revise according to that peer review, and then submit a revision with an explanation of the changes they made. I do this so I'm not wasting time going between two files to make sure changes were actually made, though I'm sure plenty of students make stuff up and submit the same draft.

Then, after they turn in their revision, they schedule a meeting with me via our Canvas Calendar. I hate Canvas, and every other LLM, as much as anybody else, but this function is amazing and I wish I could make it work through other applications. I set up 30-minute office hour blocks that students can sign up for. Once they sign up, that block disappears for other students, so you don't have to worry about double-booking. I have these appointments available from 10am to 5pm Monday through Friday, except for the times when class is meeting or other things are scheduled. I can fit in around one hundred meeting blocks in a two-week span.

And here's the thing that makes this all work: I don't even look at their paper until they are in the room with me. That way, if they miss their appointment, which a few will do, I don't waste any energy prepping feedback that doesn't happen. It takes me ten to fifteen minutes to read through their 4,000 words (give or take), and then we spend the next fifteen to twenty talking about what went well and what needs work. For the first draft, we mainly talk about developmental things: form, argument, organization, sourcing. For the second draft we talk about sentence-level mechanics... cleaning up passive voice, etc. I'll highlight the sentences and areas we talk about in Canvas so they can be reminded later of what we talked about, but I don't put any written feedback in there. They are responsible for taking down the feedback in ways that work for them.

Because they are in the room with me, they can ask questions for clarification, and I can try to gauge if they are understanding the feedback I am giving them. This is also a great way to find and correct uses of Generative AI.¹ If I see passages that sound like GenAI, or have sources that don't exist, or are mis-used, I can simply ask the student to tell me more about the source or the argument. If they've used GenAI, they generally can't. If they haven't used it, they'll be able to talk about the argument beyond what they've written. Usually if I press them, they will admit to using GenAI, which removes the necessity to try to prove it through other means, which means I don't have to be as much of a cop as I might otherwise.

¹ I forbid any use of Generative AI in this course, giving students basically one warning and then reporting them for an Honor Code violation related to using unauthorized aid and plagiarism.

The Results

The first thing I realized doing this is that it is a heavy cognitive lift to closely read and give significant and useful feedback in person in thirty minutes. I genuinely have a hard time doing it on any day that I teach, even if I can only fit a few appointments in on those days. And on my non-teaching days, doing these meetings every thirty minutes for seven hours is incredibly tiring. By the end of the two weeks of feedback I am exhausted. While the written feedback was also exhausting, and this version lasts only two weeks instead of three, it is quite concentrated. I don't have the option to take a break if I need it, unless I have scheduled one previously.

The second thing I realized is that with the student in the room, there is a real, affective, connection with them and their work. This is not something that can be quantified, but the connection seems to provide a real boost to the student's confidence that I am doing my best to help them improve and not judge them for missing some predetermined mark. Students often come into the first meeting nervous, worried I'll tell them their project or writing is bad, and leave excited to continue because they can see the path from where they are to where they want to be. I believe that this is one of the reasons that the most common written comment in my year-end evaluations is that the students feel like I really care about them and their success. Fostering this relationship is, genuinely, why I do this job, and this feedback method has been the most effective way of nurturing that relationship that I have been able to do. Even if this method had no other benefits, this would make it worthwhile for me.

That said, there are other benefits. A big one is that students are much less likely to try to sneak Generative AI into their papers if they'll have to sit in the room while I read it, and potentially get called out on it (risking an honor code violation along with a failing grade in their final course and being unable to graduate). I've only had two instances of Generative AI in the two years I've been doing this method, and both were used only for formatting bibliographies (and I tell them to use Zotero... it's literally easier than GenAI). Now, is it possible a student has slipped other instances past me? Sure. Is it possible I'm bad at spotting Gen AI use? Maybe. But the stories I hear from my colleagues about rampant obvious use make me think it's a problem I have mostly solved. Because reading through these papers multiple times, while talking with their authors, I can hear their voice and read that voice in the room at the same time. Having those conversations helps me match their thoughts to their writing in ways that just reading cannot.

I do not have good data on whether this switch affected my student evaluations. The n of a couple years of finding my feet and a couple years of doing it this way is just too little to find any statistical significance. Also, my grading rubrics have shifted over the years, so generalization of student performance compared to the actions I discuss here are basically impossible. What I can say is that in my end-of-semester reflections, about a quarter of the students mention the in-person meetings as being valuable to them. I do not prompt them to write about the meetings (though perhaps I should), and they still bring it up, many imploring me to keep it. I do have a couple who will write and say they would have preferred written feedback, but very few. Almost all will say they felt their writing improved over the course of the semester, though how much of that is just a result of multiple drafts of editing, and how much is the in-person feedback is hard to say.

The Research

To go along with this feedback process, I use an alternative grading system, sometimes called un-grading (Stommel, 2024) or specifications grading (Leslie & Lundblom, 2020). Most assignments in the course are pass/fail. If the student gives the assignment a good-faith effort and completes it, they get full points. If the work is insufficient, I will generally give them half points and ask them to revise and resubmit. If they do and there is improvement, I'll update their grade to full points. I do this to remove the "did I do this in the way that the professor will give me an A?" stress that can negatively impact students' motivation to learn (Klapp, 2015). It also removes a lot (but... sigh... not all) of the grade-grubbing and late-semester negotiation that can plague teachers as we are finishing up our courses.

Furthermore, and I'm sure this will personally offend many of you, I have no penalty for late work. Some of my assignments do eventually close and can no longer be turned in, such as old drafts when new drafts come due. In those instances, I let my students do extra credit projects of their own choosing to make up the points. For what it's worth, very very few of them take me up on this option. I had two people do so this year because they were remote, which made their participation in other parts of the course difficult.

Bridget Smyser presented a paper here at ASEE in 2024 about her study about lenient deadlines, showing that students who turned things in late often scored worse on those assignments, but that other corollary benefits (such as accessibility for disabled students,

mental health benefits, and stress reductions) likely outweighed any score decrease (Smyser, 2024). I would also note that a worse grade on an assignment is still better than a 0. Smyser and others have spoken often about how “rigor” is often framed as tight, legalistic, rules in class, difficult assignments, and a good curve of grades (whether or not they are officially curved), how these framings are damaging to student performance and mental health, and how flexible deadlines are overall helpful to student participation, success, and engagement (Becker, 2016; Miller et al., 2019; Patton, 2000; Smyser, 2024). Nearly every one of my students in their end-of year reflections mention how much the flexible deadlines helped them during their final, stressful, semesters (though some admitted it did lead them to procrastinate more).

Through specs grading, flexible deadlines, and direct feedback, I focus the students on the task of improving their craft. They worry less about having to get it right the first time, at a specific time, and instead can let that shitty first draft be shitty. They can focus on actually sitting with their writing, working through it, learning how to edit, polish, and improve over time. Every piece of pedagogical writing I have found talks about the benefits of teaching students how to self-edit rather than showing them the correct way to put a sentence together via line edits (Robb et al., 1986; Hull, 1987; Fitzgerald, 1987, just to name a few from a single year). Because we all know that if we give those significant line edits, it takes us ages, and they just change what we ask them to with little reflection on *why*. And, at least in my experience, they often fell back into the same patterns in their subsequent writing, meaning they hadn’t learned the ostensible lesson, and then I would have to do the line editing all over again.

Limitations

These methods will not work for everybody. The reasons for this can be multiple. There could be personal reasons, institutional reasons, course design reasons, and more.

One main limitation is that this does not scale for classes much larger than mine. I could possibly add another 50 students and maintain this process without needing to drop a draft. In that instance, the feedback would take three weeks for each draft, which is a long time to have zero bandwidth for anything else (see next point). For larger sections, training our TA’s – should we be lucky enough to have them – could add variation in the feedback the students receive. If we don’t have them, it might just not be logistically possible.

Another related limitation is the scheduling and exhaustion angle. I basically cannot do anything but the feedback during those weeks. I have no bandwidth for anything else. No conferences, no panel discussions, no other service requirements. I even make sure not to schedule my D&D games during those weeks because I don't have the energy or time to prepare. With the take-home written feedback model, I could have as many meetings as I wanted during the week because I would be doing the work at home. This is a tradeoff I am happy to make, but if you have a significant amount of other obligations, setting those aside for two weeks might not be possible.

Also, if you don't read quickly, this is not a system that will work for you. I know that, as academics, we are expected to be fast and skilled readers, but that is not always the case. Many of us are dyslexic, or just take longer to work through a piece. We might not feel comfortable giving feedback without the opportunity to sit and ruminate on what we have read and really get our thoughts in order.

There is also, and this might be the deadliest drawback for this method, the problem of the detail of the feedback we can give in the time allotted. You will just not have enough time to go line-by-line and give feedback on every detail. You'll be able to get big things; organizational issues, repeated writing issues like passive voice or verb tense agreement, missing or misunderstood sources, argument issues like logical fallacies or skipping evidence, cherry-picked sources, etc. Getting into the more advanced skills—transitions, rhetorical flourishes like repetition or metaphor, pacing and rhythm of sentences—is often just impossible unless you're working with upper-level students who major in a field where writing is a main skill. I work with engineers. Getting them to drop the passive voice constructions they've been trained for three or more years to use is hard enough. There is simply a level of fidelity I cannot reach through this method.

That said, I posit this may also be an advantage. Because you can't point out every instance, and can only give broad-strokes feedback, the student is then forced to find the instances of the problems you mention throughout and, if they take the task on in good faith, are likely to internalize how to see those errors and fix them, improving their skill in self-editing. Perhaps such a hope or expectation is pollyannish.

Conclusion

In conclusion, this is a feedback format that has worked excellently for me. While I am extremely busy during the weeks I am giving the feedback, I am only busy during business hours. My students have, so far, been more likely to take on the feedback I give them, and genuinely find and fix the issues I tell them about. This is not universal, of course, but it is much more prevalent than it was when I did written feedback. I have seen significant improvement in the overall quality of papers submitted in the final, graded, portion of my course, though so far the n is low for the experiment, and the variations with which I have graded things across the years make direct comparison difficult, if not impossible.

That said, I take less time overall to give feedback, that feedback feels more impactful in the room, I have significantly less instances of Generative AI uses than other instructors, and once I am done with meetings for the day I go home and play videogames. Or watch a ball game, or read a book. Once I go home, I have zero stress about grading, because it's all contained within the working hours of the weekdays. My weekends are free, I have fewer grade-grubbers than most (though some of that has to do with being a 6-foot white guy with a bass voice and grey in my beard... I was built in a lab to be taken seriously, whether or not I deserve it), and my students report that the in-person feedback was worthwhile to them in their end-of-year reviews. I've never seen a student say the same about written feedback.

With the caveats listed in the Limitations section above excepted, I can wholeheartedly recommend this method to reducing grading stress and improving student uptake of feedback, with several ancillary benefits.

Bibliography

- Becker, K. (2016). *Death to deadlines: A 21st-century look at the use of deadlines and late penalties in programming assignments*. Western Canadian Conference on Computing Education 2006.
- Fitzgerald, J. (1987). Research on Revision in Writing. *Review of Educational Research*, 57(4), 481–506. JSTOR. <https://doi.org/10.2307/1170433>
- Hull, G. (1987). The Editing Process in Writing: A Performance Study of More Skilled and Less Skilled College Writers. *Research in the Teaching of English*, 21(1), 8–29. JSTOR.
- Klapp, A. (2015). Does grading affect educational attainment? A longitudinal study. *Assessment in Education: Principles, Policy & Practice*, 22(3), 302–323.
- Leslie, P., & Lundblom, E. (2020). *Specifications grading: What it is, and lessons learned*. 41(04), 298–309.
- Miller, L. A., Asarta, C. J., & Schmidt, J. R. (2019). Completion deadlines, adaptive learning assignments, and student performance. *Journal of Education for Business*, 94(3), 185–194.
- Patton, M. A. (2000). The importance of being flexible with assignment deadlines. *Higher Education in Europe*, 25(3), 417–423.
- Robb, T., Ross, S., & Shortreed, I. (1986). Salience of Feedback on Error and Its Effect on EFL Writing Quality. *TESOL Quarterly*, 20(1), 83–95. JSTOR. <https://doi.org/10.2307/3586390>
- Smyser, B. M. (2024). *Relaxed Deadlines: Do They Provide an Unfair Advantage?* 2024 ASEE Annual Conference & Exposition.
- Stommel, J. (2024). Ungrading: An introduction. *Pedagogy*, 24(3), 327–340.