

BAAL 2026

Reframing Teachers as Co-Designers of AI-Mediated Writing Pedagogy: Structuring Feedback Through P Versus NP

Jerry Talandis Jr.
University of
Toyama

Theron Muller
Waseda University

Aston University, Birmingham, UK



Editorial note: Filler, false starts, accidental repetition, and obvious transcription errors have been removed or corrected. The speakers' sequence, substance, examples, humor, and spontaneous Q&A have been preserved.

Presentation Transcript

Jerry

Hi, everyone. Good afternoon. I hope you had a good lunch. Thank you very much for coming. My name is Jerry Talandis Jr. I teach English at the University of Toyama on the west coast of Japan.

Theron

And I'm Theron Muller from Waseda University's Faculty of Human Sciences. Thanks for being here today.

Jerry

We really appreciate it. We're going to talk today about a way of framing feedback when working with generative AI in L2 academic writing. It's based on a three-year Japanese government grant. We're currently in year three. It's about feedback literacy and AI ethics. Just really quickly about the project: it's rather complex and involved, but essentially it's about developing supportive tools for L2 learners.

One side of the project is adapting an already existing app called Auto-Peer, which is designed to promote feedback literacy, for learners in a Japanese EFL context. We're also developing a foundational-level writing textbook that incorporates generative AI into a traditional process-writing workflow.

Last year at this conference in Glasgow, our colleague Phil McCarthy joined us to present our progress on developing Auto-Peer. This year, we're not going to talk about the book per se; we're going to focus on what we learned through designing the book's workflow.

We're going to cover these points today, but we're not going to talk about the book specifically. If you are interested in that, we'll have a link at the end to another presentation we made about it. Really quickly, the student data we'll show in this presentation was obtained with informed consent in accordance with the policies of the University of Toyama.

Just a bit of brief context. First of all, educational technologies are not neutral. This is an idea that is foundational to what we're trying to do. Technology is deeply entangled in the way we teach and learn, and technology tools make pedagogical assumptions about what these things are. They focus attention on what is easy, visible, and actionable for students and teachers to do in the classroom.

Therefore, teacher agency plays a really important role in how these technologies are used. Tools provide various affordances, or ways of doing things—how do I say it?—unrealized possibilities. It's up to the teacher to make choices and decisions within their context to provide good learning opportunities for their students.

However, with the advent of generative AI, the possibilities of what teachers can do have greatly expanded. All of a sudden, within the writing process, AI can perform all kinds of roles at the same time. That has the effect of redistributing the cognitive workload of what students are able to do. At the same time, it's really important to remember that these technologies are not neutral themselves. The way they're designed steers learners toward different types of writing and interaction.

Teachers as co-designers is another key point we'd like to make. If teacher agency and judgment are so important, one way to develop them is through this idea of co-design: teachers as active co-designers who actually shape the learning experiences, either at the design stage of an application, if possible, but more commonly in the classroom.

We're looking at it this way: When does AI enter the writing process? What does it do? How much help should it provide? What should the learner actually be doing at any point? This type of conscious, mindful, active co-design is really important in the age we're in right now. That leaves us, though, with a practical question: How do you actually do that? It's kind of overwhelming. What kind of help should GenAI provide, and when? Theron, could you help answer that, please?

Theron

All right, give me the clicker. I have the power now.

As Jerry said, we've been working on this textbook that incorporates generative AI. What we don't want to do in today's presentation is go into the nuts and bolts of what the prompts look like. We did a presentation on that last year, and if you've been having insomnia because you've had jet lag like we have, there's a link to that presentation at the end. You can watch that, and if that doesn't work, there's a link to another presentation that will definitely cure it.

What we're going to talk about is the conceptual framework we've been using to think about how to design the textbook to answer two different flavors of writing problem. I don't think we're saying anything controversial here. We're borrowing a metaphor from computer science that I'll get into a little bit later.

Basically, in the writing process, there are two different flavors of problems that students may face. One is a bounded, verifiable problem: there's a clear answer to it, which we're referring to as P. I'll give you a definition in a moment. That includes spelling, punctuation, and grammatical form. There were tools

that helped students address these before generative AI came out—things like Grammarly were already addressing those easy-to-fix problems.

Then you have more open-ended, judgment-dependent things, like developing an idea, choosing evidence, and improving organization. The metaphor we're using there is NP. One of the design concepts we had as we put this material together was that, if there are two different kinds of problems students are facing, we should try to address them by creating two different solutions.

Here are the definitions of P and NP. The important thing is that this is a heuristic we're using to conceptualize them. We're not actually computer scientists. Polynomial-time, or P, problems can be solved relatively easily and have a straightforward solution. NP, or nondeterministic polynomial-time, problems take a lot of effort to solve, but once you've found an answer, it's relatively easy to confirm whether you found the correct answer.

What we're thinking about here is that problems with more straightforward solutions exist on a kind of cline in contrast to problems with less straightforward solutions. If you're a teacher, I think it's uncontroversial to say there was a problem with the organization of a student's essay. Then you look at two different essays and say, 'Oh, this one's better organized than that one.' But going from the disorganized essay and thinking about how to make it organized is a more cognitively difficult task than looking at two essays and deciding which one is better organized.

The point of this slide is that we acknowledge this is a continuum. These aren't binaries. There are things that are more straightforward and things that are less straightforward to address. There's going to be a test later. We're going to give you blank pieces of paper, and you'll be expected to fill all these in. Jerry's been working very hard to make some beautiful slides.

Basically, with straightforward problems, we're trying not to do what Grammarly does and fix the problems for students. We're trying to place the problem in front of the student and say, 'This is a problem,' then get the student to think about it and try to address it, so it's the student's cognition that resolves the problem. For the less straightforward problems, the idea is to give students options and have them think about which option is better for their own writing.

That leads to this question: What's happening when learners work within this design? I wish we had a bunch of quantitative data for you, but we're going to give you two anecdotes. Like I said, there's a test later.

This is the overall design we went through. The thing to keep in mind is that the students are guided through writing with something called flipped prompting, where the AI asks them questions and they answer. Then there's P—trying to fix some simple errors—the students reflect, there's traditional peer feedback, and finally—sorry, I broke my arm; I shouldn't be waving it around so much—at the end of the process, the generative AI suggests changes to them. At that point, they think about what changes were made and whether they agree with those changes. They're not getting the changes and then submitting them as a final assignment.

These are the two anecdotes we're going to walk through. We have one student who followed the intended trajectory—that will be pretty quick—and another student who was an edge case and didn't follow the trajectory we intended. That student's learning path, or interaction path, ended up informing the design we implemented.

This is the typical student, or the student who did what we expected. On the left is the original version of the assignment, and on the right is the revised version. You can see the nice red circle there. 'Think himself' was changed to 'think for himself.' I think it's pretty uncontroversial to say that's a simple problem to fix. In the revised version, it was fixed. Thank you, Jerry.

The whole point is that the AI didn't fix it for him. The AI asked him to think about a solution, and then he thought about the solution and fixed it.

This is the second, more interesting student. Their task was to write 150 to 200 words. In the original version, they had less than a third of what they were asked to write. In the final version, they were close to the minimum of 150. The big mystery to us was how they ended up with the extra text in the rectangular box.

In a lot of the dialogue, there's an implicit accusation that the generative AI did it for them, or that they asked the generative AI to do it for them and they cheated. Thankfully, the way Jerry set up his research, we could audit the interactions the students had and see what actually happened in the process. This isn't the whole interaction; it's just to give you an idea of what happened.

The student says, 'I have to write this paragraph, 150 to 200 words, so I'd like to expand it.' Then the AI starts asking questions, which is what we want. We don't want the AI to write for them; we want the AI to get them to think about how to expand the story. That looks good.

The issue we found when we looked in more detail is here: the middle column is what the student is writing, and the right column is what ends up in the student's text. The student gives relatively minimal answers, and the AI fills them in. That isn't really what we want. We want the AI to give the student what they initially wrote and then give the student an opportunity to expand it. Seeing how the student interacted with AI led us to revisit and revise the design to address these kinds of issues, which we'll get into next.

The next part of the progression I showed earlier—where I said there's going to be a quiz—is for the students to get peer feedback. This is the peer writing to the student, saying, 'Reduce the number of pronouns because it was hard to tell who was doing what.' I think a number of writing teachers here would say, 'Actually, you want to have more pronouns because that's the way English is expected to be written.' But one point to keep in mind is that, in almost all cases in the writing literature, people are supposed to be writing for an audience.

In this case, these students are writing for an audience of their peers. Their peers are all at about the same level, so the use of pronouns makes it more difficult for them to read the text. I'm not saying it's right or wrong. I'm saying that when you're thinking about which reader you're trying to address, you have to think about how you structure the text for that reader. In our conceptualization, this is the kind of problem that doesn't have a clear answer.

You can use 'Hana' to refer to the character in the text, or you can use 'her' as the pronoun. While some people would say the pronoun version is preferable, both are grammatically viable choices. One isn't clearly right and the other clearly wrong. When you think about the readability of the text and its intended audience, you might choose differently depending on that audience.

The student presents the text to the AI, and the AI recommends changing it to 'her.' The student says, 'No, I don't want to do that. I don't want to use "her." I want to refer to the person by name.' That is

based directly on the feedback they got from their peer about readability, and it connects to the idea of the learner retaining authority over the formulation of their final text.

The next thing we're going to move to is how meaningful learner work can be preserved when GenAI enters the process. Here we'll address how we redesigned the unit. I'm going to pass it over to Jerry.

Jerry

Over a year and a half of iterative development, I think I can boil down what we learned to three main points. Number one: if you want your students to use AI, don't let them use AI at the beginning. They need to do their own thinking first.

We rely on that old, traditional free-writing task. Give them 10 minutes to think about it and write freely. Then they go into the AI having already established their own thinking. They're in the driver's seat rather than being a passenger in their own writing process.

Another lesson, as we showed earlier, is that the AI needs to be programmed to elicit, not to automatically jump in and fill gaps. That's really important. The third, very important point is to tell the AI what you don't want very clearly. That's as important as telling it what you do want.

We certainly didn't really get that at the beginning, as you can see from some examples here. We had to focus on tightening the wording of the prompts with more specific language. We added new guardrails, and, as I mentioned earlier, we sometimes changed the structure of the prompts. For example, instead of having AI enter at the planning stage, we began with focused free writing. These are some examples of how we altered our design.

If you apply this to the P–NP heuristic, when you're dealing with P-like problems, you want to constrain the AI's role to these types of tasks so the learner can do their own thinking. This is like programming AI to be a personal tutor, in effect. With NP-like problems, the aim is to program the AI so it proposes ideas, offers alternatives, or lays out versions of the student's writing and the AI revision side by side in a table.

Then we ask students to reflect on and justify the language they incorporate from the AI. We take the AI's writing and use it as an object of study and comparison in order to develop critical thinking. That was a key insight we gained. But what does that mean for teachers and researchers?

Theron

These are some of the implications we wanted to go through. Keeping in mind the advice from this morning's plenary not to tell teachers what they should or shouldn't do, the solution we found that works for us in this context is for the teacher to act as an orchestrator.

Rather than passing the technology over to the students and saying, 'You figure it out,' the teacher stays involved in the process and the design process, and thinks about how cognitive responsibility is delegated. We worked on the prompts to make the guardrails much more explicit and kept an audit trail to see how the learners actually interacted with the technology.

That's where we are with the three-year project we're concluding now. We've applied for a five-year grant, and we'll find out in February whether we get it. Our thinking is moving from a much more localized question—how students interact with AI, and how to design for students to interact with AI—to thinking about it as a kind of human–AI learning ecology. Thank you.

We began with: How does GenAI participate in L2 writing? The answer is: it depends. Now we're thinking about how to orchestrate this process. How do you develop a human–AI learning ecology that takes the generative AI in the process into account? The idea is not to consider AI as a property of the tool, but to consider it as part of a larger ecology and ask how to integrate that tool into the larger ecology. And that's it.

As promised, those are the QR codes. If you're having some bouts of insomnia, this is our talk last year at BAAL.

Jerry

This one goes into the actual prompts and how we designed them. This one is about the textbook in more detail, if you're interested. So, yes, just check it out.

If you have any questions, we have five minutes or so. Any comments or questions? Do any of you teach writing, or are you struggling with AI and all that?

Q&A

Questioner 1

Thank you. This is a fascinating presentation. I'm not a computer programmer, but I have an argumentative question. When you were talking about how a student would put in minimal input and ChatGPT would produce this expanded output, it reminded me of how we talk in second-language acquisition about recasts: a student says something, and then the teacher repeats it but expands it.

I guess what I'm asking is, if we think of it in those terms, to what extent does ChatGPT become an aid, at what point does it become something that is doing the work for them, and at what point does it become neither because the student doesn't have uptake of the recast?

Jerry

That's an interesting question.

Theron

I think the issue for us was that, according to the design principles, at that stage we wanted the students to be producing the language. We didn't want the computer to be producing the language. There's a later stage where the student takes the language they produced and says to the computer, 'Here's what I have. How do you make it better?'

From that point, the assignment is: What did the computer produce? What did you produce? What's the same? What's different? What do you think about what's the same? What do you think about what's different? It becomes a reflexive task rather than, 'Oh, here's my final assignment.'

As you could see from some of the scans we had, one of Jerry's strong beliefs is that, even if they do tell the computer, 'Write it for me,' he has them fill out a worksheet by hand. They're handing in a piece of paper on which they've actually written out text. Even if the computer produces text, they're the ones—if it says something like 'floccinaucinihilipilification,' they're the ones who get to figure out how to write that out by hand. So they're at least thinking about it to a certain extent; they're at least doing that motor skill of writing it out.

Jerry

I would also add that it's an interesting idea to take that back to what we were saying about teachers being involved in the whole process. If that's what you're noticing, then you can make that a choice of what to focus on in a lesson. That's what I was hinting at about the possibility space being so vast. You have all this opportunity to intentionally look at something like that as a recast. You can turn anything into a positive with the right approach.

Theron

Just to follow up on what Jerry said, this is an example of a design solution. This is not the example of the design solution. We're not trying to say to anybody, 'This is what you should do.' We're keeping in mind the advice we received from the plenary this morning. We're saying this is an example of what we, as teachers, found worked for us.

Jerry

We're just trying to figure it out. What a mess, the situation we're in. I hope some of this has been helpful in helping you think about it.

Basically, if you're involved in the process and aware of what your students are doing, there are ways of working with AI where you can turn things around and use it in a beneficial way. We're still exploring that, and I hope we can all keep working at it and figure it out eventually. Anybody else?

Questioner 2

Thanks for the talk. I'm just curious—I'm sorry if you talked about this last year—but how good is this program, or how good is GenAI, at determining P- versus NP-like problems? You did say it's a continuum. Do you find it's generally pretty good at distinguishing the two?

Theron

That is an excellent question, and that is what we're hoping to address in the grant proposal we submitted. We don't have a clear answer to that; that's the future research question. We want to know that too, right? That's the £50,000 question. The answer is whether we get the £50,000.

Questioner 2

I reckon it's pretty good, from my experience, but that'll be interesting. Thank you.

Theron

Maybe one more, if anybody has one. Okay, we've got the next presenter waiting, so we'll get out of the way here. Thank you very much.