

AI-Laoshi Will See You Now

How Artificial Intelligence is Transforming Chinese Language Learning

Jeremiah Jenne
David Moser

Jeremiah Jenne: Hello and welcome to another edition of Barbarians at the Gate. This is your host, Jeremiah Jenne, broadcasting from Geneva, Switzerland. And with me, my co-host, David Moser. How are you doing, David?

David: Pretty good. A little frustrating. I can't get my camera to work, but then my eyes aren't working too well either because I just came back from New York and I have blurred vision maybe from jet lag. So you don't want to see me. I look like a mess.

Jeremiah Jenne: Last time I was in New York too, I had blurred vision, but that has a lot to do with O'Flaherty's in Midtown. So yeah, I could totally sympathize with that. It's good to have you back in Beijing. It is the start of another academic year. And it seems like every year or every few years, there's a new innovation out there in the world of technology that causes teachers like us to throw our hands into the air collectively and proclaim the end of education as we know it and possibly the end of civilization. Whether it was the dawn of television back in the 40s and 50s to the internet, then Wikipedia, and of course now it's AI, AI everywhere.

So David, let's kick this off. I want to talk today a little bit about what AI means for teachers, what it means for students, what it means for study abroad and also a little bit too, because you and I are both researchers, we work on a lot of academic projects in our own scholarly lives. How does AI affect our research? How does it affect our writing? How does it affect our thinking? What have you seen in terms of AI in the classroom, whether with study abroad students or with Chinese students?

David: I was amazed at how fast China was able to develop its own GPT type of AI with DeepSeek. Most of my students are aware of AI, aware of ChatGPT and these things. I just met with some teachers the other day and said students are downloading these things and there are lots of them and they're starting to use them. I think there's a lot of consternation, a lot of fear involved, a lot of, there's a sense of looming chaos. But I think the problem now is what do we think of these tools? Are they the end of the world, people getting lazy, there's no way to actually run proper classrooms, or can they actually be used as a tool and a tool that would revolutionize the way we teach?

I have long been an advocate of students taking advantage of the digital technology to teach themselves as much as they can independently. I think you will agree with me, I'm sure, that the biggest game changer of the 21st century has been Pleco. That app absolutely revolutionized the way we learn Chinese and the way we read Chinese. And a lot of people thought, well, is this cheating or should we let people? I don't think there was very quickly people adopted it. And I think it's universal now. There were some competitors, but Pleco seems to be the best. And it is a godsend for overseas students, they're all using it.

But this AI is a different matter because it's not just a super smart dictionary that doesn't use paper and you don't have to worry about radicals. You probably realize when you talk to young students learning Chinese, they don't even know what a radical is. They say, isn't that someone marching in a parade with a sign? They don't even have any idea. They don't need to know what a radical is, right?

Jeremiah Jenne: They missed the experience of standing in front of like an entire bookshelf with a multi-volume like *Hanyu Da Cidian* 汉语大词典 and sitting there on your palm, counting out stroke order so that you could look it up in one of the indexes. And to be honest, I have to admit that was the case when I first started learning Chinese, but then in my academic career, it's coincided with the rise of Pleco and Wenlin and some of the early, as he said, paperless dictionaries, I think, is a good way of describing.

David: Right. That's right. Yeah. So now we have a device that's so powerful, it's not only going to revolutionize how students learn, but about how teachers teach. We've long ago had programs that would allow teachers to make their own class notes, class tests or exams, or study sheets easily, all formatted and everything. Most of the teachers that I've been working with didn't use those very much. But the next wave of teachers are going to be people who grew up in the digital world. They're going to use these things when they teach Chinese because it's a breeze.

When ChatGPT first came out, I experimented with coming out with test sheets that would, I could just give it a command and see how well it would do. So one of the tests that my Chinese students are concerned with is something called the TEM4 test, test of English mastery or something. And it's something that all undergraduates have to take at the end of the second or third semester in order to get qualifications to be able to teach English. So it's a very important test. I just gave ChatGPT the instructions. "Devise for me a CLOZE test," if you know what that is, C-L-O-Z-E. It's the sort of multiple answer test. You have a blank and you choose A, B, C, or D.

And I just said, "Print out a close test for the TEM-4 at the appropriate level of English difficulty, choose a topic, print out all of the questions with the blanks." And it was just incredible. Within seconds, it came up with something that was actually amazingly good. The topic was electric vehicles. Very good topic. And it came up with some questions. It had a passage and then it had the blanks that you had to fill in for the appropriate vocabulary word. And it was formatted and it was formatted exactly like the TEM4. I said, this is astounding. This is amazing. So I checked it and everything seemed correct. But then I noticed the correct answer for the number one question was A. The correct answer for the number two question was A. The correct answer for the number three question was A. And it was my first realization that it is at once the smartest thing you've ever seen and also the dumbest thing you've ever seen.

Now I tried that again recently and it did not do that. It did not make that mistake. So within a year or two, that was corrected, right? So now we have this very powerful and getting even more powerful. I think we're at ChatGPT-5. DeepSeek is also pretty amazing, but they're different. They have a different style, but the functions are basically all the same.

So the thing that's happening right now, Jeremiah, is for dedicated students, I think we should throw away the paper and start letting students independently or at least uniquely develop with the teacher their own curriculum or their own sets of skills that they need. This has always been a problem really with Chinese, as you know, because we were always stuck with these boring handouts or texts that were out of date as soon as they were in our hands. They were boring and uninteresting. And so we depended a lot on the teachers to bring all that to life.

There's a risk of being dangerous because it's going to make people lazy, people are going to cheat. I mean, there's like two types of people. One type of person says, gee, I don't need to really learn Chinese anymore. I just have all these things I can use. And there will be technology where you can stand and talk to another person with earphones and just get a simultaneous translation. So people say, great, we don't have to go through all this hassle. I can just communicate using all this technology. But the other more dedicated and I think more reasonable students or the more serious students will say, great, this is a way that I can learn what I need to know and what I need to say and what I want to say and express to myself and use to take what I need from the technology to quickly get up to a level where I need to be.

I've always had this philosophy of when you're starting, concentrate on the vocabulary and the sentences that you need to use every day. So if you're a vegetarian, you need to know how to say that. If you're Jewish, you might want to mention that. If you have a peanut allergy, you might want to mention that. But it was a little bit hard for teachers to be able to help every student with that. Nowadays, I think teachers need to just totally revise the way they teach.

The study abroad programs are going to have to, I mean, they should sit down and just totally revamp the curriculum, totally rethink it. Because, don't you think that after a semester or so, students are going to think, why am I taking this Chinese class? It's irrelevant to what I have today. And people are going to go off and learn Chinese on their own. I've already met some several young whippersnappers who really want to come to China and they've told me. They said the classes, the teachers are wonderful, but the class material is worthless.

Jeremiah Jenne: I have two thoughts on that. The first one is that during my teaching career, and of course in study abroad, you often have faculty. It was my experience, and this may have been limited to the programs where I was teaching, but the area studies faculty, the content faculty was often, and maybe it was due to the subject matters or whatever, was often more willing to adopt new technology than the language teachers. And part of that is because many of the language teachers were incredible teachers who were fabulously well-trained in a particular method of language acquisition. And to deviate from that, I think was difficult.

And that led to a phenomenon, which I'm sure you've heard. We probably said this on the podcast a few times. You have students who come back to you, either at the end of the year or during the year or several years later. And they say, yeah, it's a great program. But I'll tell you, I learned way more Chinese, fill in the blank, playing ultimate frisbee, hanging out at the bar, rock climbing, whatever, than I ever really did sitting in the classroom. And I always thought to myself, well, that's not entirely true because you're getting the foundation in the classroom that then you used in the bar or on the, playing ultimate frisbee. I get the point in that, just as you said, you tend to acquire language in the areas that are most meaningful to the student, to me, to

you. And what ends up happening is of course we end up with a memory bank of vocabulary that we download at appropriate times.

I'm in a restaurant, here's the restaurant list. I'm talking about history, here's the history list. I'm talking about jazz, here's the jazz list. Acquiring and building those lists is often the part of language acquisition, at least for me, that was the most exciting. And here's where I think it dovetails with AI. It is that you can have conversations with AI. You can, as you said, develop textbooks based upon your own interests in a very short amount of time.

And it is quite eerie. And maybe this is part of the problem for those, if you're really worried about the long-term ethics of AI, how quickly many of these, whether it's ChatGPT or Claude or DeepSeek, mirroring the personality and knowing what you're looking for. And so you can train the model pretty easily to start having conversations about stuff that you or I want to have conversations about. And I've seen, I have a practical real world example of this in our own household. My wife, we live in Geneva. My wife who had studied for a couple years in France as a student and still speaks French better than I do, but for many weeks has an, I don't know what you call it, a ChatGPT ami that she has conversations with in French.

And they have a great time and it actually replaced her online French teacher that was living somewhere in South America on a beach and supporting himself with online teaching. And she was pretty impressed with how well it improved her French because it was talking about stuff that she wanted to talk about. And in language acquisition, language teaching, as you know, you keep them interested and that's how you keep them interested and you keep them asking questions about how do I say this?

David: I must say I was a little bit stunned. It's an eerie feeling talking to the computer. One of the things that improved massively over the years is intonation. We had text to speech and various things, but not only is the language, in this case Chinese, perfect, but the intonation is also very human, almost indistinguishable from a real human being. Some of the things that ChatGPT can do, it will do anything you want it to really, but you can ask it for role playing dialogues. The computer plays the role of a taxi driver or a boyfriend or whatever. And you can speak as much Chinese or English as you want. They accommodate that or it accommodates that.

And you can have a conversation at whatever level you want. And if you don't know the word, you can switch to English and say, how do you say sofa in Chinese? And they'll either answer in Chinese or English, whatever your preference is. If the AI is speaking too quickly, you can request that it slow down, and it will do that. It's kind of a boring interaction in some ways in that it's just incessantly congratulatory and positive. It will always say, great, good question, really good.

Jeremiah Jenne: But you can change that too. You can ask it to be a hard ass and it will be.

David: That's right, but it's not really what it's designed for. And it's a bit humorless, I think. Anyway, we're getting too picky. It's already amazing that you can do what it does. You can also ask for a weekly exercise regime where you do a certain amount of tones and a certain amount of vocabulary, a certain amount of sentence structure. So something that for anyone who's serious

and really wants to learn Chinese, this is something that's really a game changer. Because as you say, what your students reported to you is you're just not likely to get anything like that in a classroom in the traditional structure that it used to have.

I've found it useful because no matter how long we've been speaking Chinese, we still make errors. And very often, the better you are, the less you really realize the errors you're making because if you're fully communicative and the other person can always understand what you're getting at, then you sort of no longer have the feedback that it takes to show that you're saying something in a wrong way because no one cares what mistakes you're making. They just want to understand you. So when you reach a level that you and I are at, I think you really need people who are still pointing out your errors. And that's something that you wouldn't get with just an ordinary talk with someone else, because they're not going to be concerned with your language problems. They just want to understand what you're saying. So as I say, this is a godsend for people who seriously want to learn Chinese, but it's also a disaster for the huge number of people that won't even try anymore.

Jeremiah Jenne: Yeah, I think that idea about correcting errors is an important one, particularly if the primary person you speak with is the person that you're also in a long-term relationship with. Because, I mean, there's two parts of that. One, you do get accustomed to the way someone speaks a second language. So you start to just kind of not, you don't hear the mistakes anymore because you know how they speak, but also there's objective hazards to pointing out errors especially in your wife's Chinese, especially when she speaks multiple languages better than I do.

David: I feel your pain. I know what you're talking about.

Jeremiah Jenne: Yes, it's one thing to kind of use it to practice like how, like I'm going to a cafe in France, help me to practice how I should say this so I don't sound like I'm a complete idiot. But what happens when we then take this to another level? So for example, you work with academic texts or literary texts. I often work with classical Chinese texts and this is where, when I first started using Pleco or Wenlin, where the problem started because those had dictionaries that were classical dictionaries or even some of the big ones like *Hanyu Da Cidian* 汉语大词典 multi-volume epic dictionaries. But when you're translating Qing dynasty, Song dynasty, classical texts, there were limitations to just, if one...

For those of you who've ever studied classical Chinese, when one character can mean 43 different things, that's a little difficult for any electronic dictionary to process. And what I'm surprised by is how much AI gets right in translating even somewhat difficult classical Chinese and also what it tends to get wrong. It's like the last, when they talk about delivery like logistics they always talk about the last hundred meters, right? It's the challenge. I think and this is true, I think with AI generally speaking, it helps to get you most of the way there in a shorter amount of time than it ordinarily would take you but it doesn't always bring you home or if it does you got to be really careful about the route it takes because that's where the errors creep in. And I think that's why being a serious student of Chinese or any language is important because you need to be able to read over, be able to at least read over the translation and understand how they got to what they got. Like, okay, I see how you translated this this way, because I can see the line here

and I know now where the error is because that one character is always a little tricky, that kind of thing. And if you don't have that foundation, the AI can actually cause greater problems. And if I may, this is why in daily life and research, in all things, I generally think of AI not as artificial intelligence, but as the artificial intern. And it tends to do a job. And I say this with due respect to all interns out there. I was an intern once myself. And as interns will always be, they do a great job of helping, but you have to check the work.

Sometimes the helping actually creates more work or problems than the help might otherwise solve. Have you tried to use AI for, if you will, serious work or for academic work?

David: I would say right now, if you use just not even the AI programs, but just the translation apps, Google Translate and DeepL and these sorts of things, very often you discover that it's a matter of input. Right now, I think most of those apps, especially the ones that are developed for Chinese, they actually have lots of, for example, Tang poetry. The standard translations or the standard interpretations of Tang and Song poetry are just there to be pulled out. It's not really a translation exercise. They recognize that poem, they'll give you a standard translation of it.

And it's possible they've got so much memory and computer power now that they could really take all of the translations of the *Lunyu* 论语 and the *Zhuangzi* 庄子 and everything and poetry and get them all in the database so that you could actually get some very good standard translations. Of course, the other thing is if it's *wenyan* 文言, if it's classical Chinese, if you look at translations done by human beings, even Chinese native-speaking human beings can differ in their interpretations. This is an aspect of classical Chinese that's unavoidable. There's a lot of ambiguity, and part of the beauty of it is that you do have to pull the meaning out in semantics from a very terse and information-stingy text. So I think this is only a matter of time when the artificial intern knows more than you do, or at least has more database than you can hold in your pitiful brain. And you'll find that I think it's going to be a boon to people not only doing standard modern Mandarin, but also classical language.

For the kind of research I do, it's psycholinguistics. It works pretty well. I'm very suspicious sometimes, but I've checked. I now have five different AI programs that I use just to...

Jeremiah Jenne: Which ones if I may ask?

David: So that's Grok, DeepSeek, the one that used to be called Gemini, the Google one, and of course, ChatGPT. And there's some other ones that aren't very good. I use just basically those four or five. But it's very useful because, and I think it's getting more and more useful because they continue to upgrade and enter more information. I think these are going to be tools that are extremely powerful and are going to be an aid to sinology. It already is. A lot of the boring work like doing statistics, it's a breeze now. I just finished editing a book that's coming out in the fall. And all of the indexes are a nightmare because there's these different formats that you have to use, Chicago or AP or something like that. And you could just put all these references in ChatGPT and just say, make an index in alphabetical order with all of these references. And within three seconds, they have it. And it's things like that as a scholar that save you an inordinate amount of time. The sorts of drudge work that you have to do, especially with Chinese

when you're going back and forth between English and Chinese and pinyin. These computers are going to make that a breeze but with a little caveat, it still can make mistakes. And when you're just putting your faith in the program and throwing all that data at it, it's very easy to get lazy. What happened with Robert F. Kennedy Jr., remember he came out with this paper that was about American eating habits. And the index had a lot of imaginary articles that didn't exist, which were lazily using ChatGPT or something, and it was hallucinating articles that didn't exist. So I think it's going to be the same thing in academia. You have to be very careful to make sure that the computer, that the AI is not making mistakes.

Jeremiah Jenne: One last question, David, since it's only natural for us as scholars, as researchers, as academics to use these tools in our own work and integrate them in our own work, is it unrealistic, unfair to expect students to not use these tools in their own work? And if they are using it, what are some ways that we or even you have employed to change perhaps how we evaluate or teach students that either allows for a certain amount of use of AI or somehow prevents use of AI to take shortcuts that could inhibit the learning process?

David: Yeah, well, this is a good question. It really is. It's hard to know what kinds of abilities and functions that we should keep and ones that we can now mercifully discard. And that's a really big question.

I think that it really is reasonable to demand that all students learning Chinese actually practice translating or writing in Chinese or speaking without the use of AI to make them at least minimally or at least adequately competent to create sentences and to do translations just with your brain.

I think if people get so lazy that they never do the translation on their own, I think that's dangerous. I think that's not a good thing. I think you need to force them to at least go through the rigmarole of looking at the text and try to figure it out without the machine just giving you the answer.

But we have to adapt to this digital age and we have to maximize our resources, right? You take a math class now and you don't need a slide rule. That used to be *jiben gong* 基本功 back in the old days. You had to know how to run a slide rule. Now we have calculators. We don't use them. So the same thing is happening in the field of studying Mandarin.

Jeremiah Jenne: David, I feel like this is a subject that in the next months, years, we're going to come back to in various forms, in terms of how AI affects different parts of our job, academia, because I do think like in all industries, like in all fields, the rise of AI, how to integrate it into what we do is obviously going to be a topic of some discussion for study abroad. And there's a number of ways that we can explore this in the future, for example, how does it affect teaching students about cultural interactions? How does it affect how students process information of what they're seeing? Do they ask the person next to them or do they ask ChatGPT? All of these things, I think are part of an evolving landscape for education in general and for study abroad.

I really appreciate you allowing us to kind of kickstart that conversation. And I hope those of you who are listening will tune in for future episodes as we come back to this topic.

David: Yeah, let me just add one thing. I think that the goal we should be striving for is to let these electronic devices aid in interacting with other human beings. The goal should be to improve your Chinese reading, writing, speaking so that you can function with other human beings communicatively and not make this mistake of getting students just addicted to their cell phone and the screen and the computer and spend all their time interacting with it. That would be a disaster. So I think it's very important to keep that goal in mind. The goal is to improve your interaction with other human beings, not to get addicted to another computer device. Don't you think?

Jeremiah Jenne: I totally agree. And I also think we should be sure to teach all of our students to say, "I want you to go and find the one they call Sarah Connor and have her terminated" in multiple languages. So they know when the AI has finally taken that drastic last step. Thank you all for joining us today. We appreciate it. We'll be back in two weeks. And as they say, cue the, I guess they're not AI, it's more of a sound clip, but cue the drums.