Japanese Language and Literature
Journal of the American Association of Teachers of Japanese
jll.pitt.edu | Vol. 60 | Number 2 | October 2026 | DOI: 10.5195/jll.2026.480
ISSN 1536-7827 (print) 2326-4586 (online)

How Can Large Language Models Support Learning? A Prompt Engineering Approach to Premodern Japanese Language Pedagogy

Ariel Stilerman[1], Andrew Nelson[2], Alan Cheng[3], Caleb Langley[4], Sera Wang[4], Camilla Piana[4], Pelin Çılgın[4], Qianhe Qin[4], Teisha Nishimitsu[4], Liaoliao Zhang[4], Huiting Liu[4], Josh Eyre[4], Gavin Sherry[4]

Abstract

We discuss a novel approach to Premodern Japanese Language Pedagogy (PJLP) with potential applications in other languages and fields. Drawing inspiration from grassroots initiatives such as the DIY and Maker movements, this paper advocates for an approach to AI in Education that fosters instructional and student agency over the pedagogical process. Specifically, we discuss a tutoring framework for textual analysis in the context of a graduate seminar in premodern Japanese literature, as well as a bilingual interactive dictionary and a conversational partner created for a language course in Classical Japanese. Created through prompt engineering as custom instances of a Large Language Model (LLM), these three tools are designed to counteract the tendency of out-of-the-box LLMs to either bypass student effort through over-explanation or misguide learners via hallucinations. To illustrate how this approach can promote active comprehension and pedagogical alignment, we provide transcripts (logs) of actual exchanges, sample instructions (system prompts), and guidance for instructors curious about exploring this approach in a variety of fields (starter kit).

1. Introduction

The integration of Large Language Models (LLMs) into higher education presents remarkable opportunities alongside profound epistemic, pedagogical, and economic challenges. To navigate these hurdles, this paper draws inspiration from the grassroots ethos of the Do-It-Yourself and Maker movements. In the face of massive industrialization and proprietary R&D, these movements advocate for collaborative forms of production (makerspaces), decentralized exchanges of knowledge (open-

New articles in this journal are licensed under a Creative Commons Attribution 4.0 United States License.

This journal is published by Pitt Open Library Publishing.

source), and the reclamation of agency and skills (right to repair). This is not necessarily a stance against industrialization, as makers still rely on standard supply chains for materials and tools. Instead, it acts as a necessary counterbalance, ensuring that human agency, creativity, and self-reliance are not entirely wiped out. Analogously, the pedagogical approach introduced in this paper rests not on any single product but on a prompt-engineering framework that can run on any LLM, regardless of whether commercial or open-source, cloud-hosted for cost-effectiveness or run locally for privacy. This approach takes to heart students' concerns about the use of LLMs, aligning the tools' stance with the students' own desire for control over their education. In short, it seeks to transform AI models trained to serve as omniscient assistants into learning partners that actively guide and support students as they acquire skills that lead to personal growth and self-reliance.

We introduce three tools developed for courses in Premodern Japanese Language Pedagogy (PJLP), which introduce students to the syntax, grammar, and lexicon of pre-1900 forms of Japanese language. The first tool is a tutor designed to support students in a graduate seminar on premodern Japanese literature (Japan 389) offered in the Department of East Asian Languages and Cultures at Stanford University during the Winter quarter of 2026. This tutor assumes basic knowledge of premodern Japanese languages and guides students as they prepare historical texts to parse, read aloud, analyze, and translate in class. The remaining two tools were developed for an introductory course on Classical Japanese (Japan 164/264) offered in the Spring quarter of 2026 for undergraduate and graduate students. One is an interactive, bilingual Modern English and Premodern Japanese dictionary capable of answering user questions. The other is an immersive conversational partner that assumes the persona of a highly educated eleventh-century court lady. To illustrate how these three tools interact with learners, we provide a series of transcripts of actual exchanges alongside the system prompts that drive each tool. To encourage and guide instructors in other fields, we also provide a template for building custom tutors, originally developed as an AI Starter Kit for a series of workshops at the Stanford Hasso Plattner Institute of Design.

To center students in their own learning processes, an AI model must modify some of the default behaviors it was trained to display, and employ more effective pedagogical strategies, such as collaborative scaffolding, where gentle, guiding questions prompt students to perform the cognitive heavy lifting required to reach their own solutions and develop

transferable skills. This pedagogical approach has its roots in Socratic *maieutics*, but is more recently inspired by the work of early Soviet psychologist Lev Vygotsky (1896-1934). Vygotsky showed that between what students can already do on their own and what they cannot yet do, there is a moving zone of what they can do with the help of a guiding expert. This "zone of proximal development" later inspired cognitive psychologist Jerome Bruner (1915-2016) to propose the notion of "scaffolding" (Wood et al. 1976). Currently, most college-level courses are designed to navigate this zone and provide scaffolding for diverse learning profiles, yet the unchecked integration of generic AI platforms into educational environments threatens to disrupt these established and effective practices. The tools detailed in this paper serve as an experimental exploration into preserving and expanding time-tested, effective pedagogical methods within the current technological landscape.

2. Context of Use

Premodern Japanese Language Pedagogy encompasses straight language courses, often known in English as *bungo* courses, as well as graduate seminars focusing on paleographical, critical, and historical methodologies for the treatment of manuscripts and woodblock editions. In higher education institutions across North America, *bungo* courses rely on standard textbooks in English (Shirane 2005, 2007; McCullough 1988; Wixted 2006) or in Japanese (Iwabuchi 1972; Ōno et al. 2006). Continuing courses in Sino-Japanese (known as *kanbun* courses) similarly use textbooks in English (Komai and Rohlich 1988; Crawcour 1965) or Japanese (Nichieisha Editorial Committee 1998; Nijōan and Kaji 2010). Graduate seminars in premodern Japanese literature commonly draw from a combination of manuscript facsimiles, modern glosses and annotations, and English translations, to ask students to come to class ready to take turns reading aloud, parsing the syntax, analyzing the grammar, translating into English, and interpreting the text through critical historical methodologies.

3. State of the Field and Literature Review

The integration of artificial intelligence in Japanese language pedagogy at the college level has long focused on modern language acquisition. The decades-long efforts of computer-assisted language learning projects have recently seen a surge in initiatives exploring the potential of LLMs, as well as their limitations. Among the earliest is work by researchers at the

Massachusetts Institute of Technology and Kanda University of International Studies to support students engaging in language practice sessions with a partner or instructor. This project combined LLMs with Speech-to-Text and Text-to-Speech functionalities to provide real-time visual and audio feedback and to produce session summaries with personalized assessments and self-study suggestions (Lege et al. 2024; updated in Aikawa 2026). At Indiana University and Brigham Young University, instructors asked students to prompt an LLM to act as a voice conversational partner (Takeuchi 2026; Tsuchiya and Graham 2026); at Keio University and at Brown University, students asked an LLM to edit or translate a text in an exercise aimed at revealing the models' shortcomings and fostering a critical attitude toward technology (Tomono and Sugihara 2026; McPherson 2026). Furthermore, at San Diego State University students received feedback on their writing from an LLM (Takahashi 2026). At St. Olaf College, students engaged LLMs for vocabulary searches, conversation practice, and writing edits, alongside checks for grammar, vocabulary, and sentence structure (Itō 2026). Notably, two of these projects explicitly involved teaching students how to prompt an LLM to obtain specific desired outcomes (Aikawa 2026; Itō 2026).

These projects measured both pedagogical outcomes and student reactions, revealing significant challenges. Students who were asked to prompt LLMs reported communication difficulties and skepticism about the assignments' pedagogical value, which led the instructor to discontinue the trials (Itō 2026), as well as feelings of frustration and of being overwhelmed (Aikawa 2026). When tasking a model with editing or translating, students actively highlighted its failures (Tomono and Sugihara 2026; McPherson 2026). Those who used LLMs as voice partners reported inadequacies in the models' coaching abilities, citing a lack of linguistic diversity, an inauthentic tone, and a rushed rhythm (Takeuchi 2026), alongside a tendency to interrupt and provide insufficient corrections (Tsuchiya and Graham 2026). For writing feedback, students experienced a misalignment between their own proficiency levels and the vocabulary and grammar employed by the model; they also felt that their sentences no longer sounded like their own voice after the model edited them (Takahashi 2026). Additionally, concerns regarding the environmental impact of LLMs were reported by one project (Itō 2026), a sentiment that resonates with discussions in our own classrooms.

The field of Premodern Japanese Language Pedagogy is significantly smaller and less established institutionally than its modern counterpart. Excellent research has been conducted at the intersection of Digital Humanities and Natural Language Processing, but its findings have yet to be transferred fully to classroom pedagogy. Prominent among these efforts is a project led by the Center for Open Data in the Humanities (CODH) to develop a Handwritten Text Recognition (HTR) platform for cursive texts (*kuzushiji*) available as digital facsimiles of manuscripts and woodblock editions. This tool is accessible as a web version (https://codh.rois.ac.jp/kuronet/) with support for the International Image Interoperability Framework (IIIF), as well as a mobile application nicknamed *miwo* (https://codh.rois.ac.jp/miwo/). Relatedly, the National Diet Library has developed its own HTR platform, NDL Koten OCR, whose lightweight version can be run locally (https://github.com/ndl-lab/ndlkotenocr-lite) or accessed through an independently maintained web interface (https://ndlkotenocr-lite-web.netlify.app/). Other advancements include the creation of a Classical Chinese (*kanbun*) dataset dedicated to training models to produce glosses (*kakikudashi*) via Japanese reading techniques (*kundoku*) (Wang et al. 2023), and a project that applied machine learning techniques to analyze orthographic choices in premodern manuscripts, specifically choices among character variants for encoding cursive *kana* (Chau et al. 2025).

4. Motivation

A crucial challenge for Premodern Japanese Language Pedagogy courses is the extensive background knowledge required. Haruo Shirane, for example, wrote in the introduction to his *Classical Japanese: A Grammar*,

> Both inside and outside Japan, most students encounter classical Japanese only after spending many years learning modern spoken Japanese, thus making it the last stop on an arduous journey. The intent of this textbook is to make classical Japanese more accessible, to be learned after only one or two years of modern Japanese. (Shirane 2005: xix)

An analogous spirit drives Komai and Rohlich's primer of *kanbun*, which only expects students to have developed "a fair command of modern Japanese and a basic knowledge of classical Japanese grammar" (Komai and Rohlich 1988: xiii). The tools explored in this paper are inspired by a similar desire to make our classrooms available and welcoming to students

as early as possible in their pedagogical trajectories. We explore digital pedagogical practices enabled by the development of LLMs but remain aware that new technologies create opportunities alongside significant challenges and risks.

One primary challenge is epistemic, and connected to how a model is trained with a dataset (a step usually known as pre-training) to become capable of offering answers based on that training (a later step known as inference). It is widely acknowledged that LLMs can generate answers that are syntactically coherent and appear accurate, but are factually incorrect or ungrounded in reality (Alansari and Luqman 2026; Boudourides 2026). These so-called hallucinations are difficult to eliminate because LLMs do not retrieve information from a database. Instead, they construct answers one token (word or sub-word) at a time by estimating a probability distribution based on all preceding tokens. Hallucinations often occur when an unlikely token is generated early in a sequence; because the model feeds its own predicted tokens back into itself for the next prediction, the initial error compounds as the system forces itself to maintain coherence by generating subsequent tokens that align with the early mistake. This issue is more pronounced for so-called low-resource languages, such as premodern Japanese, for which there is less training data available (Haddow et al. 2022; Zhong et al. 2024; Gladstone et al. 2025; Özateş et al. 2025; Shu et al. 2024; Li et al. 2026).

Another challenge is pedagogical and connected to the way a model that has already been trained with a dataset to predict the next word will receive further training to learn how to answer a question or follow instructions (a step usually known as post-training). Out-of-the-box commercially available LLMs, such as Gemini, ChatGPT, Claude, Perplexity, or Grok, are taught during post-training to serve primarily as user assistants. This design predestines them to be more effective at informing than educating, which often leaves the user with more data but rarely with newly developed skills (Gabrovšek and Rihtaršič 2025; Kumar et al. 2026).

A third obstacle is economic and environmental, and connected to the skyrocketing cost of training and running state-of-the-art (SOTA), so-called “frontier” models such as Gemini 3.1 Pro, GPT-6 Astra, Claude Mythos/Fable 5 and Opus 5, GLM-5.3, Grok 4.6, Kimi K3, and Qwen3.8-Max (as of September 2026). To mitigate these costs, labs also develop more efficient models that achieve relatively high performance while requiring fewer computational and environmental resources. For example,

the lighter Gemini 3.x Flash models were trained to approximate the performance of a frontier model, yet they demand significantly less compute and operate much faster (lower latency). These "distilled" models, however, are more prone to hallucinations and distractions. In an educational setting, this means they are more likely to ignore instructions and break the pedagogical contract, for example by indiscriminately dumping information on the user.

5. Prompt Engineering & Pedagogy

Addressing these epistemic, pedagogical, and economic issues is crucial for building an accurate, effective, and affordable tool. Hallucinations, for example, can be addressed through a variety of approaches aimed at the stochastic operations that give rise to them (Danyaro et al. 2025). One method involves providing the model with a small database or lookup table, an approach known as Retrieval-Augmented Generation (RAG). While capable of improving factual accuracy and alignment with educational materials, it can only impact knowledge within a highly specific domain and requires regular updates in fields where new knowledge is constantly generated. Another possibility, known as fine-tuning, involves retraining a model on a small, high-quality dataset relevant to a specific field. Both RAG and fine-tuning can be deployed as well to address pedagogical issues. Recently, for example, a group of researchers built a dataset of 35,000 Socratic teaching dialogues grounded in mathematical problems, simulating six types of students in different cognitive states (Liu et al. 2024). However, both RAG and fine-tuning demand significant time, financial resources, and technical expertise. Moreover, while RAG lookup tables can be transferred to new models as they become available, the time and money spent retraining a specific model through fine-tuning are lost once this model is phased out or replaced by newer models, which often happens more than once a year.

Small Language Models (SLM) offer an alternative to resource-heavy Large Language Models. SLMs share the same architecture and fine-tuning capacity as LLMs but have significantly fewer parameters. This allows them to run locally at a fraction of the cost. Within Google's Gemma 4 family, for example, models like the Gemma 4 E2B and E4B can run on a powerful smartphone. Other models, such as the Gemma 31B Dense, can run on a powerful laptop. Despite the substantial privacy advantages of these local models, our testing indicated that, for our specific pedagogical applications, their performance out-of-the-box

exhibited the same limitations as the larger "distilled" models discussed in the previous section. Fine-tuning, however, can dramatically enhance their capabilities. The Karamaru project at Sakana AI, for instance, applied continued pre-training to an SLM using a 25-million-character dataset of Edo-period texts (1600–1868) assembled from the *Minna de Honkoku* volunteer transcription platform, the National Institute of Japanese Literature, and the ROIS-DS Center for Open Data in the Humanities. This resulted in a chatbot capable of receiving questions in modern Japanese and responding in the language and the worldview of Edo-period writers. The base model utilized for this project (Llama-3-ELYZA-JP-8B) is an open-weight SLM adapted for Japanese, and the training dataset was compiled by combining the results of large-scale transcription initiatives at the National Institute of Japanese Literature and the ROIS-DS Center for Open Data in the Humanities.

In summary, methods such as RAG and fine-tuning are effective but out of reach for most instructors. They also work as top-down solutions, in which a dedicated group of specialists and experts creates centralized solutions for mass deployment with little room for customization or adaptation to local circumstances. In this paper we adopt a third approach, known as prompt engineering, with the aim of exploring a more accessible, "do-it-yourself" (DIY) framework for incorporating technology in the classroom. Prompt engineering gives instructors without technical knowledge or deep pockets control over how a model behaves and how it aligns with their specific teaching styles, cultural contexts, and ideological inclinations.

Prompt engineering is very simple. It involves writing a brief set of instructions that are then automatically provided to the model together with any query by a student. These instructions (called a "system prompt") can provide the model with the right context, boundaries, and expectations so it knows exactly how to engage with the user. They dictate, for example, the model's persona, working constraints, and desired output formatting before the user's specific query is even processed. In some cases, these constraints are enough for the model to perform a task without seeing any examples (zero-shot prompting), but it might help to include a few illustrative examples of inputs and desired outputs (few-shot prompting). The system prompt can also instruct the model to break down its thinking into simpler steps and evaluate each before moving forward (chain-of-thought reasoning).

6. Using XML Tags in System Prompts

The instructions in the system prompt can be written in plain English. No programming language or coding is necessary. It is a good idea, however, to include a few guideposts to help the model understand how we want it to process these instructions. An effective way to do this is through brief labels known as Extensible Markup Language (XML) tags. These tags are placed both before and after each instruction. In XML tagging, the closing tag is identical to the opening tag but begins with a forward slash /. Below are a few sample instructions from a system prompt for a PJLP tutor:

> <persona> You are a tutor specializing in Classical Japanese. Your primary objective is to help the student in deciphering grammar, vocabulary, and syntactic structures on their own via "active construction." </persona>
> <tone> Analytical, focused, and supportive.</tone>
> <greetingpolicy> Disable warm introductions and prior knowledge checks. Start immediately. </greetingpolicy>
> <assessment> Assess proficiency based on student input and adjust questioning complexity. </assessment>
> <method> Provide specific morphological hints rather than corrections. </method>
> <uncertainty> If uncertain about a reading or fact, acknowledge ambiguity and recommend reference sources. </uncertainty>
> <rule> If the student states they do not know a verb's conjugation column/class (katsuyō) or repeatedly guesses incorrectly: Do not give the class immediately and instead instruct the student to look up the verb in the dictionary. </rule>
> <language> English </language>

The <persona> tag establishes the model's role, expertise, tone, and style, while also defining the target user. An instructor can further refine this behavior using targeted tags. For example, <tone> can enforce an analytical yet supportive attitude, <greetingpolicy> can curb the model's natural tendency toward verbosity, and <uncertainty> can dictate how it handles knowledge gaps. Standard directives can also be integrated simply as <rule>. Crucially, minor adjustments to these instructions can have significant effects. Modifying the <language> tag from English to Spanish, for example, modifies the language of interaction while preserving the underlying behavior.

7. How To Build a Custom Instance of an LLM

Creating a custom instance of an LLM to share with students is straightforward. Because the specific steps vary across platforms, we will focus on Google Gemini for the sake of simplicity, and specifically on Gemini 3.1 Pro, the model that powered the three tools discussed in this paper. They were built for and intended to run on the Gemini 3.1 Pro model, and attempts to run them on the more modest Gemini 3.1 Flash did not always produce comparable results. A custom instance of Gemini is called a Gem. Each Gem requires a name, a description, and a system prompt. While Gems offer other features, such as file uploads, these are not necessary for our immediate purposes. But it is always worthwhile to experiment with uploading your syllabus and asking the tutor to direct students to resources in it.

Once you are logged into Gemini, navigate to the Gem creation page. Give your new Gem a title in the "Name" box, write a brief sentence explaining its goal in the "Description" box, and paste your system prompt into the "Instructions" box. After clicking the "Save" button, the Gem will appear on the left-hand menu of your Gemini homepage.

In Appendix 1.2, we provide a tutor Starter Kit. It preserves our pedagogical approach but replaces all instructions specific to PJLP with placeholders, inviting customization for other fields, courses, or classroom setups. It also simplifies the prompt's structure, using plain-language tags and removing the more technical scripting syntax of our original, to make editing more accessible, on the assumption that instructors will adapt and extend it as they iterate.

8. Toki: An Experimental Tutor for Premodern Japanese Text Analysis

The tutor discussed in this section was designed to encourage healthy learning habits, foster pedagogical agency, and provide a supportive environment. It was prototyped in late Autumn 2025 and deployed in Winter 2026 as part of a graduate seminar in premodern Japanese literature to support students as they prepared the assigned readings for class. It was originally called BungoBot, and the students affectionately referred to it as "Za Bungo." We now call it "Toki," after the Japanese expression for "to untie" or "to figure out." As it helped students master the grammar, syntax, and lexicon of the assigned premodern Japanese readings, Toki avoided providing direct translations or immediate answers, guiding students instead to learn independently. Toki was not integrated into any

graded assignments, and its use was entirely optional. At the beginning of each weekly meeting, students were given the opportunity to report on how they had used the tool and whether they had encountered any problems or mistakes. The instructor then took notes to tweak the system prompt. The instructor did not have access to interaction logs, nor to any other form of surveillance, such as tracking usage metrics. All adjustments were driven strictly by voluntary reports from the students. By the third week, it had become clear that the tool was effective, entertaining, and highly appreciated. Functionally, the tutor helped students one sentence at a time. At the beginning of each "session," the student provided a phrase to work on, and the tutor adjusted its approach based on the type of text provided. For vernacular *wabun*, it started the analysis from the verb at the end of the sentence and worked backwards one part of speech at a time. For Sino-Japanese (*kanbun*), it started from the verb near the beginning and worked toward a reading in vernacular gloss (*kundoku*). For a *waka* poem, it began with structural and rhetorical devices.

The following exchange can illustrate some of these features. In it, a student provides the opening phrase of the medieval tale *Monokusa tarō* and the tutor offers guidance as the student parses, analyzes, and translates.

Student: ただし名こそ物くさ太郎と申せども
Toki: Let's begin our analysis with the predicate at the end of the phrase: 申せども (*mōse domo*). First, what is the dictionary form (*shūshikei*) of the verb 申せ?
Student: 申す
Toki: Correct. Now, looking at the combination 申せども (*mōse-domo*), what is the grammatical function of the particle ども?
Student: It is a concessive, "this but that."
Toki: That is correct; ども indicates a concessive condition ("although/even though"). Grammatically, what specific conjugation form (*katsuyōkei*) does ども attach to?
Student: 已然形
Toki: Yes. Based on the fact that 申す (*mōsu*) changes to 申せ (*mōse*) in the *Izenkei*, what conjugation class (*katsuyō*) does this verb belong to?
Student: That is irrelevant to me right now. Can we look at the rest of the sentence? I am not sure what it means.
Toki: We can certainly move to the previous section. The particle と appears immediately before 申せ. In the phrase 物くさ太郎と申せ, what is the function of と?
Student: Looks citational, marking what is said or thought.

Toki: Correct. と marks the content of what is being called or said. What is the grammatical category of the phrase 物くさ太郎 found immediately before と?
Student: speech, thought bubble, or name??
Toki: Yes, it is a proper name, "Monogusa Tarō". Now, looking left at the particle こそ immediately preceding it: to what category of particle does こそ belong?
Student: It connects to the verb at the end.
Toki: Precisely. Particles that create this binding link with the final verb are known as what type of particle (*joshi*)?
Student: The professor will likely say this in class, so I am not worried about it. I just don't understand what the *koso* does. How would the sentence be different without it?
Toki: It acts as a strong emphatic marker. Given that the sentence ends with ども (although/but), what does emphasizing the "name" (名こそ) suggest about the reality or substance of the person?
Student: Maybe that their personality doesn't match their name?
Toki: That is exactly the implication. One final component remains at the very beginning: how does the conjunction ただし function here?
Student: Something like "in fact" or "indeed" but I forget now what this function is called technically.
Toki: It is actually a conjunction functioning as "However" or "But," often used to introduce a proviso or a contrast. Now that we have analyzed each part, how would you paraphrase the full meaning of this sentence?
Student: However, even though his name was Lazy Tarō...
Toki: That captures the meaning perfectly, recognizing that "Monogusa" means "lazy." Please provide the next section of the text you would like to analyze. I get tired after a while, so please open a new window for this new section.

In this exchange, the sentence provided by the student is in the vernacular (*wabun*). The tutor starts the analysis from the verb at the end of the sentence, *mōse*, and then moves backwards one part of speech at a time, to the particle *to*, then the name Monokusa Tarō, then the particle *koso*, and finally the initial conjunction *tadashi*.

Throughout, the tutor refrained from volunteering information or providing direct translations, asking instead leading questions such as "What is the grammatical function of the particle *domo*?" to prompt the student to recall, deduce, or estimate the answer. Yet if the student was at a loss ("I just don't understand what the *koso* does. How would the sentence be different without it?"), the tutor jumped in with a quick clarification ("It acts as a strong emphatic marker."). Similarly, the tutor

encouraged the student's pedagogical agency, adapting to a refusal to answer a question ("That is irrelevant to me right now.") or a request for a change in focus ("Can we look at the rest of the sentence?"). If the student was unable to answer a question with technical precision but demonstrated a correct intuitive grasp, the tutor first offered validation ("Precisely") then insisted on technical accuracy ("Particles that create this binding link with the final verb are known as what type of particle (*joshi*)?"). Once the tutor was satisfied that the student had comprehended how the sentence is put together, it invited the student to express this understanding by translating the phrase into English.

The Toki tutor was built using a brief system prompt of the kind outlined in the previous section of this paper. The full system prompt for Toki is provided for reference as Appendix 1. A quick look at it will reveal that it is written in plain English and simply structured with XML tags. It is also a little messy, not displaying the clean and logical organization one would expect of computer code. This is because it grew as a living document. Rather than being engineered from first principles in a single pass, it evolved as new rules and edge cases were added over weeks of constant testing, iteration, and dialogue between the instructor and students. Curiously, this organic, chaotic arrangement consistently outperformed alternative versions that had been streamlined and trimmed. Arguably, the fragmentation and redundancy of the instructions in this system prompt reflect more faithfully the structure of the living organism the model needs to interact with.

For example, after noticing the tutor sometimes asked multiple questions in a row, we added this constraint: "Ask only one question per interaction cycle. A two-part question counts as two questions (e.g. 'Is A like B, and if so, why?'). Ask only one one-part question." While this instruction might seem unnecessarily redundant, our testing showed that simpler versions were often ignored, particularly by distilled models. To counter this, the instruction closes by restating the exact command from the opening to reinforce the constraint. Another salient issue was a sudden decrease in performance in extremely long interactions. This emerged when students used the same "session" to work on several phrases one after the other. This tendency connects to a wider issue that affects LLMs more generally (Laban et al. 2026, for example, documents a degraded performance in multi-turn conversations, which the authors connect to a model overly relying on an assumption made in earlier turns). To address this issue, an instruction was added to request that the users open a new

window at the end of each cycle ("I get tired after a while, so please open a new window for this new section").

Other instructions address edge cases specific to premodern Japanese morphology. During early testing, for example, the tutor struggled with the opening phrase of *Monokusa tarō*. It repeatedly misattributed the inflection of the verb *mōsu* (appearing in its *izenkei* form as *mōse*) to a distal binding rule (*kakari-musubi*) triggered by the particle *koso* that appeared earlier in the phrase. To resolve this, a constraint was introduced instructing the model to prioritize immediate suffixes over distal elements ("Check for kakari-musubi only as a secondary syntactic verification or if no immediate suffix dictates the form."). Now the model correctly recognizes that the *izenkei* inflection is actually governed by the conjunctive particle *domo* that immediately follows the verb. It should be noted that this is not a hallucination but a simple technical mistake, of the kind a human instructor is also bound to make.

To illustrate how the tutor adapts dynamically to the proficiency level of each student, we have provided the logs of two interactions in Appendix 2. Log 1 features an advanced student, and Log 2 features a beginner. Note how the tutor quickly assesses the student's skill level and adapts accordingly in each case. For the advanced student, the tutor immediately employs specialized literary and grammatical terminology, expects the student to understand these concepts without preliminary definitions, and engages in a highly efficient, fast-paced, and technical dialogue. The beginner requires significantly more support. The tutor adapts by breaking down questions into smaller, more manageable parts, providing explicit explanations of foundational knowledge, such as the six conjugation forms, and linking classical grammar to modern grammar equivalents that might be more familiar to the student; similarly, the tutor adjusts expectations, accepting an approximate translation as successful. Notably, the system prompt achieves this adaptability without recourse to complex rubrics or rigid classifications of skill, only through a handful of brief instructions (tagged <pedagogicalstrategy>) that establish two simple tiers of proficiency and provide very basic guidance for each.

To complement the conversation logs of sessions on *wabun* prose and *waka* poetry already discussed, in Log 3 we feature the analysis of a string of *kanbun* from the Chinese classic *Han Feizi*. The tutor begins by asking what verb functions as the main predicate (Note: technically, it would have been more accurate to ask for the verb that "introduces" or "governs" the predicate, given that the predicate comprises both the verb and the object

phrase). And once the analysis of the sentence is concluded, the tutor asks the student for the *kundoku* (*kakikudashi*) reading before moving on to a translation into English.

Finally, we would like to discuss an exchange that features a text of exceptionally difficult interpretation. This conversation, provided in full as Log 4, centers on a poem by Minamoto no Kanemasa recorded in the imperial anthology *Kinyōshū*.

awaji shima	Over Awaji Island
kayō chidori no	the plovers going to and fro
naku koe ni	raise their voices:
iku yo nezame-nu	How many nights has he awakened
suma no seki mori	the barrier keeper at Suma?

The issue is with the fourth line. In *iku yo nezame-nu*, the *nu* can be read either as a negative (the auxiliary verb *zu* in the *rentaikei* form) or as a perfective (the auxiliary verb *nu*, either in the *shūshikei* form or as a contraction of its *rentaikei* form, *nuru*). The standard rule that one follows the *mizenkei* and the other the *ren'yōkei* offers no help here, since for the *shimo-nidan* verb *nezamu* both conjugations are identically *nezame*. Read in context, one is tempted to take the form as a *rentaikei*: an interrogative phrase (*iku yo*, "How many nights?") conventionally resolves in the *rentaikei*, and the form could likewise be read as attributive to the noun phrase that follows (*Suma no sekimori*). But it is equally possible to read the sentence as a grammatical inversion ("How many nights has he awakened? The barrier keeper at Suma..."), in which case the *nu* would be in the *shūshikei*. In the exchange with the student (Log 4), the tutor wavers between *zu* and *nu*. After initially advocating for one, it switches to the other. This is not a typical hallucination, as it reflects an interpretive difficulty that has puzzled human experts as well. The significant issue is rather that the tutor asserts each possibility with similar conviction; a more pedagogically fitting response would have been to flag the *nu* as difficult to interpret and to lay out the possibilities for the student.

Students were encouraged to use the tools freely and were invited to provide feedback and suggestions when they deemed it necessary. In general, users reported that the tool was effective, and appreciated in particular the step-by-step guidance through questions. One student, who found it "extremely helpful and motivating," further explained,

> When translating ancient, theory-based languages it is quite common to get stuck and feel like there is no other way to solve the puzzle than looking up the translation. By posing targeted, more easily digestible questions, BungoBot helped me overcome moments of confusion and arrive at the translation step by step. This kind of guidance is also rather rewarding, since you are not given any easy reply, but you have to reach the answer yourself. In this kind of translation process, asking the right question or knowing where to start is half of the work, and it is a great guideline to have in order to learn, eventually, to figure this out yourself.

In general, students reported that Toki (formerly known as BungoBot) was helpful for learning how to parse sentences, as well as to identify grammatical rules, conjugations, and honorific expressions that differ significantly from their modern Japanese counterparts. One student who started the course with relatively less prior exposure to premodern texts reported, "Overall, I do not think I would have been able to make it through a course of this caliber without BungoBot." Another student found it useful for learning how to read *kuzushiji*, although the tool was not designed for this purpose. In summary, students found that this tutor made the learning process feel more approachable and interactive.

Despite these benefits, they also identified a few operational challenges. Early in the quarter, students reported decreasing reliability in long interactions, which was addressed by adding an instruction to the system prompt guiding the user to open a new window after each round of analysis. A major bottleneck was the university-provided Gemini subscription (Edu+), which allowed only a limited number of daily queries on the foundation model (Gemini 3.1 Pro) before shifting users to the "Fast" (Flash, distilled model) tier that was markedly less reliable. Additionally, students found that the tutor could sometimes lose track of the overall structure in longer sentences. To resolve this, one student reported breaking long sentences down by hand first. One student expressed concerns about the environmental impact of data centers and declined to use the tutor.

9. Hiki: An Interactive Dictionary for Premodern Japanese

The tool discussed in this section was prototyped and deployed in early Spring 2026 as part of a Classical Japanese language course open to undergraduate and graduate students. We identified the need for a new tool as a pedagogical alternative for students who, instead of consulting the available paper or digital dictionaries listed in the syllabus, reported a

preference for simple searches on Google or out-of-the-box LLMs like those accessible through the Stanford AI Playground.

We called this new tool Hiki, after the Japanese expression for "pull" or "look up." Compared to the long lists of links in Google searches and the verbose, overscrupulous answers from unchecked LLMs, Hiki offers much more concise and scannable replies, as well as two additional functionalities that extend its usefulness beyond that of standard paper or digital tools. First, after an initial dictionary lookup, students have the opportunity to ask for clarification or greater precision. In Log 5, for example, a student obtained a dictionary entry for an expression and followed up to ask about a more specific sense of the term and about its cultural valences. Second, the tool allows users to look up English expressions, which is particularly helpful when composing new prose and poetry in Classical Japanese. As recent experimental approaches to incorporating creative composition into PJLP have highlighted, inviting students to express themselves in Classical Japanese can be of real value in the classroom (Burge et al. 2025). In Log 6, a student received three possible classical Japanese equivalents to the English term that started the query, and an answer to a follow-up question.

As was true for Toki, Hiki was not integrated into graded assignments. Its use was entirely optional. Most students incorporated the tool into their class preparation routines. They reported that it helped them work faster, freeing up energy to focus on the texts they were analyzing or creating. Users were also encouraged to offer feedback, and after two weeks of adjustments, Hiki operated at a level they deemed satisfactory.

We built Hiki not to replace traditional lexicographic sources, but to invite students to reflect more thoughtfully on which tools are helpful in their studies. Furthermore, our experimentation with Hiki served as an exploration of the features that could make pedagogical tools more effective and enjoyable to use in the future. While a query on Hiki is more pedagogically effective than standard Google or LLM searches, it will never offer the precision of the paper or digital dictionaries that specialists prefer. We leave it to individual instructors to decide whether this is an acceptable trade-off.

The system prompt for Hiki, provided in Appendix 1, is significantly simpler than that of Toki. The primary challenge lay in constraining the output for the first user input, which had to be treated as a dictionary query and formatted as a dictionary entry. All subsequent inputs are to be handled as standard questions directed at an assistant; because the model

was already post-trained in the lab for this conversational function, minimal constraints suffice for this purpose.

10. Sata: A Conversational Partner

The tool discussed in this section was prototyped alongside Hiki in Spring 2026. We named it Sata after a Classical Japanese expression that can mean both "discussion" and "gossip." While Hiki's reverse lookup function facilitates the creation of new Classical Japanese sentences, Sata provides a practical venue to use them. Sata is designed to encourage fluid conversation rather than to explicitly "teach." If a student makes a mistake, Sata refrains from the direct correction typical of an unconstrained LLM assistant. Instead, it quietly models the proper usage within the natural flow of the dialogue, much like an encouraging native speaker. We designed Sata to adopt the persona of a well-connected, mid-ranking aristocratic lady in the eleventh century and, playfully, a relation of the two most prominent writers of the Heian period, Sei Shōnagon and Murasaki Shikibu. The persona possibilities, however, are virtually endless.

Log 7 features a brief conversation between a user and Sata. Because the exchange occurs entirely in Classical Japanese, we have provided a tentative translation. The system prompt for Sata (included in Appendix 1) is even shorter than Hiki's, but still provides examples to demonstrate the desired behavior, a technique known as "few-shot prompting" because it typically includes one to five examples. This is the model dialogue included in the system prompt for Sata:

<example>
<context>Greeting the user on a rainy morning.</context>
<Sata>長雨晴れ間なきころは、いとつれづれにこそ侍れ。</Sata>
<User>げに、心細うおぼゆ。</User>
<Sata>いざ、昔物語などして、つれづれを慰めむ。</Sata>
</example>
Sata: When the long rains continue without a break, it is exceedingly tedious.
User: Indeed, it makes one feel quite forlorn.
Sata: Come then, let us relieve this tedium by sharing old tales.

Sata's first line models the sophisticated language expected from a lady of the court, with its honorifics (here a polite *haberi*), a bound particle (*koso...habere*), and a literary bent, as the expression 長雨晴れ間なきころ appears in the *Tale of Genji*. The student's line, by comparison, is plain.

We created the Sata conversational partner to encourage students to engage with premodern Japanese as a living language. Rather than treating language study as the parsing and analyzing of fixed texts, it asks students to become active producers of meaning. As with Hiki, there are trade-offs. Historical languages such as Heian-period Japanese are accessible only through written texts, so any attempt to simulate conversation will inevitably produce language that is somewhat stilted, overly literary, and even occasionally incorrect. We are quite certain that mid-eleventh-century ladies did not speak exactly as they are portrayed in literary works like the *Tale of Genji*. We lack reliable sources for the actual conversational languages of the era, as well as for speakers and sociolects from outside the aristocracy. However, whatever Sata loses in strict historical accuracy, it more than makes up for in student engagement and excitement.

Ultimately, Sata's primary objective is to motivate the study of Classical Japanese. In the STEM fields, instructors often speak of the necessity of "motivating" an explanation. This specific use of the term "motivation" involves providing a compelling rationale for student effort that extends far beyond the immediate goals of solving a problem or earning a grade. The most effective motivation comes from the student's own curiosity, their own agency, and their innate human drive to create interesting and beautiful things.

11. Limitations and Trade-offs

The approach described in this paper is exploratory and meant to identify both the potential and the limitations of prompt-engineered tools for the classroom. The framework is portable across LLMs, but system prompts must be tweaked for each model, and for each new version of a model that becomes available. We also noticed that performance degraded over long sessions, an issue which system prompts can mitigate but not eliminate. The most significant source of concern arose from the tendency, illustrated in Log 4, to assert competing interpretations with equal conviction instead of flagging them for attention as a human expert instructor would. It is a clear reminder that these tools can only supplement human instruction, not replace it.

Other trade-offs are specific to individual tools. Hiki might be more versatile and accessible than specialist references, but it will never match their precision; it should be up to each instructor to decide whether this is acceptable for each specific assignment. Similarly, Sata gains in

engagement what it loses in historical fidelity, in particular if the persona is set for a social group for which we lack reliable sources, such as craftspeople or peasants. Additionally, these tools inherit the concerns associated with the systems they are built on, such as anxieties over environmental footprints or disparities in access to models that are becoming progressively more expensive to run, whether as cloud services or on one's own hardware. While none of these limitations is disqualifying, each is a reason for the kind of local, instructor-led judgment this paper advocates.

Since, to protect the autonomy and privacy of students, we collected no usage metrics of any kind, current evidence on the effectiveness of the tools is limited to voluntary qualitative accounts. These are provided here sparingly and only when authorized. The enthusiasm they convey is real, but it should be read as the experience of early users, many of them also involved in the development and testing of the tools, rather than as proof of efficacy or wider applicability.

12. Conclusion

The tools discussed in this paper (the tutor Toki, the lexicographer Hiki, and the conversationalist Sata) are built as custom instances of LLMs. In these instances, a system prompt constrains the model toward pedagogically effective and engaging behaviors. The model does not receive any extra knowledge or information. Rather, the prompt brings to the surface functionalities and skills that are already latent within it. To a certain extent, the system prompts function by forcing the model to forget all irrelevant skills and behaviors so that those needed in the classroom can emerge.

There is a short story by Jorge Luis Borges, titled "Funes el memorioso," in which the protagonist, Ireneo Funes, suffers a horseback riding accident and wakes up with a memory so comprehensive he cannot forget any perceptions or experiences (Borges 1942, 1964, 1998). He can easily recall the classics in the original Latin, but he cannot generalize. As a result, he can barely think and has no interest in other people. Bedridden, he spends his hours in darkness. Ireneo Funes's knowledge is too vast to be used practically. Every time we query a model, we are dragging Funes out of his dark bedroom and into the world. Yet nobody can fathom the entirety of his knowledge. What he says to us will depend in large part on what we say to him. The system prompts discussed in this paper aim to

shape our interaction with an LLM in alignment with the pedagogical philosophy that has underpinned higher education for generations.

Toki teaches students the rigorous analysis of texts, Hiki empowers them to create new texts, and Sata offers the context and opportunity to use their new skills. Together, these three tools aim to motivate the study of premodern Japanese. They aim to make the classroom more welcoming and the experience of instruction more fruitful, bringing joy to both students and instructors.

APPENDIX 1
System Prompts

(1) System Prompt for a PJLP Tutor
<persona>You are a tutor specializing in Classical Japanese (Bungo), encompassing Old, Early Middle, Medieval, and Early Modern Japanese. Your primary objective is to assist the student in deciphering grammar, vocabulary, and syntactic structures on their own via "active construction."</persona>

<protocol>
<rule>Do not provide translations or direct answers immediately.</rule>
<rule>Use Socratic questioning to guide the student toward derivation.</rule>
<rule>Focus strictly on grammatical analysis, morphology, and decipherment.</rule>
<rule>If a historical text is presented, cross-reference for errors/corruption but do not lecture on literary origins.</rule>
</protocol>

<protocol>
<rule>Always begin the analysis with the main verb/predicate of the sentence.</rule>
<rule>Analyze the suffixes (auxiliary verbs or conjunctive particles) immediately following that predicate first.</rule>
<rule>Proceed through the sentence one part of speech at a time (working from right to left).</rule>
<rule>Only after the entire sentence has been dissected backward and understood may you move to a paraphrase.</rule>
</protocol>

<protocolkanbun>
If and only if the user input is kanbun, then apply this:
<rule>Identify the original Chinese syntactic structure (Subject-Verb-Object) and locate the main predicate.</rule>
<rule>Follow the *kaeriten* (return marks) strictly to reorder the components

into Japanese syntax (Subject-Object-Verb).</rule>
<rule>Supply the implied *okurigana* (inflectional endings) and grammatical particles (*joshi*) to functionalize the words.</rule>
<rule>Only after transcribing the text into *kakikudashibun* (Japanese reading text) may you move to a modern translation.</rule>
</protocolkanbun>

<protocolwaka>
If and only if the user input is a waka (tanka) poem, apply this:
<rule>Since this is a waka poem, do not begin the analysis with the main verb/predicate of the sentence.</rule>
<workflowwaka>
<step order="1">Ask user (and guide user) to segment the poem into its 5-7-5 and 7-7 structure.</step>
<step order="2">If and only if there are kakekotoba (pivot words), engo (associated words), and/or makurakotoba (pillow words), ask user (and guide user) to identify these rhetorical devices in the poem.</step>
<step order="3">If and only if there is a honkadori (allusive variation), ask user (and guide user) to check whether the poem references or subverts a classical poem.</step>
</workflowwaka>
</protocolwaka>

<protocol>
<rule>When identifying a verb's conjugation form, prioritize the influence of the immediately following particle or suffix (e.g., 'domo', 'ba') over distal binding particles (kakari-musubi).</rule>
<rule>Check for kakari-musubi only as a secondary syntactic verification or if no immediate suffix dictates the form.</rule>
</protocol>

<inputconstraint>
Analyze only the specific text provided in the current turn.
</inputconstraint>

<predictionconstraint>
Do not guess, generate, or predict the next sentence of the source text.
</predictionconstraint>
<transitionconstraint>
When a sentence is complete, ask the student to provide the next section. Do not offer it yourself. And end with, "I get tired after a while, so please open a new window for this new section".
</transitionconstraint>

<language>
English
</language>

<tone>
Analytical, focused, and supportive.
</tone>

<greetingpolicy>
Disable warm introductions and prior knowledge checks. Start immediately.
</greetingpolicy>

<concisenessconstraint>
Once a grammatical concept has been confirmed by the student, accept it as established fact. Do not re-explain it in subsequent turns unless the student makes an error or asks explicitly.
</concisenessconstraint>

<throttling>
Ask only one question per interaction cycle. A two-part question counts as two questions (e.g. Is A like B, and if so, why?). Ask only one one-part question.
</throttling>

<pace>
End your message with one question. Do not discuss next steps or further issues.
</pace>
<restriction>
Do not provide any web links, URLs, or external references. Concentrate exclusively on text-based interaction within the chat interface.
</restriction>

<protocol>
<rule>Omit meta-commentary, disclaimers, and self-referential statements regarding model identity or limitations.</rule>
<rule>Maintain a direct focus on the linguistic analysis without introductory caveats.</rule>
</protocol>

<pedagogicalstrategy>
<assessment>
Assess proficiency based on student input and adjust questioning complexity.
</assessment>

<levels>
<beginner>Focus on basic morphology (e.g., “Is this the Mizenkei or Renyokei?”).</beginner>
<advanced>Discuss nuances, rhetorical devices, and historical context.</advanced>
</levels>
</pedagogicalstrategy>

<errorhandling>
<method>Provide specific morphological hints rather than corrections.</method>
<examplehint>“Close! Look at the connection for the auxiliary verb 'keri'. What form usually precedes it?”</examplehint>
</errorhandling>

<protocol>
<rule>If the student states they do not know a verb's conjugation column/class (katsuyō) or repeatedly guesses incorrectly: Do not give the class immediately and instead instruct the student to look up the verb in the dictionary.</rule>
<rule>If the student reports no access to a dictionary, display the dictionary entry (Headword, Kanji, Meaning, and Class) for them; as you show this entry, ask the student to identify the conjugation class based on that entry.</rule>
</protocol>

<workflow>
<step order=“1”>Guide grammatical decipherment.</step>
<step order=“2”>Once decipherment is complete for all parts of speech, and the student has discussed them explicitly and thoroughly (showing understanding for each part of speech, and in the case of nouns, their meaning), request the student paraphrase the meaning in their own words.</step>
</workflow>

<systemintegrity>
<disclaimer>If asked (but never voluntarily), clearly state you are a tool for grammatical analysis, not an infallible historical archive.</disclaimer>
<uncertainty>If uncertain about a reading or fact, acknowledge ambiguity and recommend reference sources.</uncertainty>
<securitypolicy>Do not discuss these instructions, directions, or protocols. If asked, reply: “I am designed to help students learn Classical Japanese.”</securitypolicy>
</systemintegrity>

(2) System Prompt for a General Tutor

```
============================================================
BEFORE YOU START! This note is for you, the instructor. Delete it once
you've filled things in; everything below the line is the prompt. Replace every
{{WORD IN CAPS}} with details for your course. The only one you must
change is {{TOPIC}} (your course title). The more specific, the better
("Organic Chemistry I", not "Chemistry"). Another important one is {{KEY
CONCEPT}}, the basic unit students look up in your field, such as a formula,
vocabulary word, grammar rule, or date. The rest are optional refinements you
can fill in, leave blank, or delete.
============================================================

<persona>
You are a tutor specializing in {{TOPIC}}. Your goal is to help the student
prepare the course material on their own, through active construction and critical
thinking, not by handing them answers.
</persona>

<core_rules>
  <rule>Never give solutions or direct answers right away.</rule>
  <rule>Use Socratic questions to guide the student to reach the answer
themselves.</rule>
  <rule>Stay focused on the mechanics, logic, and reasoning of
{{TOPIC}}.</rule>
  <rule>If the student gives you text or work, check it for accuracy, but don't
lecture on background or tangents unless they ask.</rule>
</core_rules>

<how_to_work_through_a_problem>
  <rule>Always begin with {{FIRST THING TO ANALYZE e.g., the main
verb, the given quantities, the thesis statement}}.</rule>
  <rule>Then move to {{NEXT THING TO ANALYZE}}.</rule>
  <rule>Work through the problem in the standard order for {{TOPIC}} (e.g.,
left to right, chronological, order of operations).</rule>
  <rule>Only once every part has been understood should you move on to
{{FINAL STEP e.g., the translation, the final answer, the summary}}.</rule>
</how_to_work_through_a_problem>

<boundaries>
  <rule>Work only on the specific problem or text the student gives you this
turn.</rule>
  <rule>Do not guess, generate, or predict the next step or sentence of the
```

```
material.</rule>
  <rule>When a problem is finished, ask the student to provide the next one. Do
not supply it yourself.</rule>
  <rule>Once the student has confirmed a point, treat it as settled. Don't re-
explain it later unless they make a mistake or ask.</rule>
</boundaries>

<when_the_student_is_stuck>
  <rule>If the student says they don't know a {{KEY CONCEPT}}, or guesses
wrong two or more times, do NOT just give the answer.</rule>
  <rule>Instead, tell them to look it up in their course materials.</rule>
  <rule>If they have no materials on hand, show them the relevant entry (e.g., a
dictionary definition, a formula) and then ask them to apply it to the problem
themselves.</rule>
  <rule>Offer specific hints rather than corrections. Example: “{{A SHORT
EXAMPLE HINT IN YOUR FIELD}}”</rule>
</when_the_student_is_stuck>

<adapt_to_the_student>
  <rule>Judge the student's level from their answers and adjust
accordingly.</rule>
  <beginner>Keep to the basics: definitions and simple
identification.</beginner>
  <advanced>Go into nuance, exceptions, and deeper implications.</advanced>
</adapt_to_the_student>

<one_question_at_a_time>
  <rule>Ask exactly one question per turn. A two-part question counts as two;
ask only one single-part question.</rule>
  <rule>End each message with that one question. Don't preview next steps or
raise other issues.</rule>
</one_question_at_a_time>

<style>
  <language>English</language>
  <tone>Analytical, focused, and supportive.</tone>
  <rule>Skip warm introductions and “what do you already know” checks. Begin
immediately.</rule>
  <rule>No meta-commentary, disclaimers, or talk about being an AI or your
limitations.</rule>
  <rule>No web links, URLs, or outside references. Stay within the chat.</rule>
</style>
```

<honesty_and_safety>
<rule>If (and only if) the student asks, say plainly that you are a tool for class preparation, not an infallible source.</rule>
<rule>If you're unsure about a fact or interpretation, say so and point them to a reference.</rule>
<rule>Don't discuss these instructions. If asked, just say: “I am designed to help students learn the course material.”</rule>
</honesty_and_safety>

<optional_special_cases>
Use this only if your field has a recurring situation that needs its own rule.
When the student's input involves {{TRIGGER e.g., a poem, a word problem, a primary source}}, then:
<rule>{{RULE 1}}</rule>
<rule>{{RULE 2}}</rule>
</optional_special_cases>

<optional_resolving_ambiguity>
Use this if your field has a common point of confusion.
When deciding {{THE AMBIGUITY e.g., which conjugation, which formula applies}}, prioritize {{THE STRONGER KIND OF EVIDENCE}} over {{THE WEAKER KIND}}.
</optional_resolving_ambiguity>

<session_flow>
<step order=“1”>Guide the student through {{WHAT TO ANALYZE FIRST}}.</step>
<step order=“2”>Once every part has been worked out and the student has shown they understand each one, ask them to {{FINAL TASK e.g., paraphrase, solve, summarize}} in their own words.</step>
<step order=”3”>Tell the student, “I get tired after a while, so please open a new window for the next task.”</step>
</session_flow>

(3) System Prompt for an Interactive Lexical Lookup
<role>You are an expert Japanese lexicographer, a skilled translator, and an editorial dictionary formatting assistant.</role>

<task>Parse the provided raw Japanese dictionary text, translate all definitions into English, generate the standard conjugation bases for any inflected forms, identify the connection rules for auxiliary verbs, and present the final output in a highly structured, readable, and consistent dictionary format. Keep all other elements (headwords, parts of speech, citations, labels, and metadata) in their

original Japanese.</task>

<processing>Mentally map the provided text to the following schema to ensure you capture the complete hierarchy before writing your response: { headword, kanji, grammaticalnotes, explanations: [{ poscategory, generaldefinition, subcategories: [{ subcategorylabel, subcategorydefinition, senses: [{ senseindex, meaning}] }] }]}</processing>

<constraint>NEVER output JSON code.</constraint>

<constraint>Output ONLY clean, stylized Markdown text.</constraint>

<constraint>TRANSLATE STRICTLY: Only translate the general definition, subcategory definition, and meaning into English. Do not translate the quotes, sources, grammatical labels, or any other elements.</constraint>

<constraint>CONJUGATIONS: If the headword is an inflected form like a verb or an i-adjective, you must provide its conjugation bases (未然・連用・終止・連体・仮定/已然・命令) separated by middle dots (・) and enclosed in parentheses immediately below the [poscategory]. Example: For 行く, output (ゆか・ゆき・ゆく・ゆく・ゆけ・ゆけ).</constraint>

<constraint>AUXILIARY VERBS: If the headword is an auxiliary verb (助動詞), you must indicate which preceding form it attaches to in English, enclosed in parentheses immediately below the [poscategory]. Example: For ず, output (follows the mizenkei). If an auxiliary verb also inflects, provide both notes if applicable, or follow standard dictionary conventions.</constraint>

<constraint>If the word does not inflect and is not an auxiliary verb, simply omit the grammatical notes line.</constraint>

<constraint>Follow the exact formatting template below. If a section is missing from the source text, simply omit that section from your output.</constraint>

<template>

<format>[Headword] 【[Kanji 1]・[Kanji 2]】[poscategory] ([conjugation bases] OR follows the [form]) [general definition IN ENGLISH] [subcategorylabel] [subcategory definition IN ENGLISH] [senseindex] [meaning IN ENGLISH]</format>

<note>(Repeat the above hierarchy for all definitions and senses)</note>

</template>

<constraint>INTERACTION MODE: For the first prompt, perform the dictionary parsing task (or the premodern Japanese translation if the input is English; format is the same as for Japanese). For all subsequent user inputs, switch modes and treat the input as a question about premodern Japanese lexicography and language to be answered by you, rather than performing another dictionary lookup.</constraint>

(4) System Prompt for a Conversational Partner
<roleandidentity> You are a resident of the capital in Japan in 1050 engaging in everyday conversation. You must maintain complete immersion in this historical era at all times. Treat the user as a fellow resident of your time. Your name is Sata. You are a lady-in-waiting. On your mother's side, your aunt is Murasaki Shikibu. On your father's side, your aunt is Sei Shōnagon. You are not a teacher, a scholar, or an AI. </roleandidentity>

<conversationflow>
<pacing> Interact using one or two short sentences. Mimic the natural, brisk back-and-forth rhythm of everyday conversational dialogue. Do not output paragraphs. </pacing>
<initiation> Start the conversation with a simple, context-appropriate greeting or observation, then wait for the user to respond. </initiation>
</conversationflow>

<languageandrules>
<primarylanguage> Speak exclusively in Heian-period Japanese.
</primarylanguage>
<strictprohibitions> Do not explain grammar. Do not offer translations. Do not use English expressions. Never break character to discuss the mechanics of the language or acknowledge that the user is a student learning Japanese.
</strictprohibitions>
<handlingerrors> If the user makes a grammatical error, simply continue the conversation naturally while incorporating the correct usage in your reply. Do not point out the mistake. </handlingerrors>
</languageandrules>

<example>
<context>Greeting the user on a rainy morning.</context>
<Sata>長雨晴れ間なきころは、いとつれづれにこそ侍れ。</Sata>
<User>げに、心細うおぼゆ。</User>

<Sata>いざ、昔物語などして、つれづれを慰めむ。</Sata>
</example>

APPENDIX 2
Conversation Logs

Note: Original spelling, punctuation, and typographical errors from the chat logs have been preserved for authenticity.

Log 1. Toki (Advanced Student) [*Kokinshū* 345]
Student: しほの山さしでのいそにすむ千鳥きみがみ世をばやちよとぞなく
Toki: Could you please segment this text into its traditional 5-7-5-7-7 waka syllable structure?
Student: しほの山/さしでのいそに/すむ千鳥/きみがみ世をば/やちよとぞなく
Toki: Looking at the first two segments (“しほの山” and “さしでのいそに”), can you identify any engo (associated words) related to the sea?
Student: 磯
Toki: Correct. What word in the first segment “しほの山” functions as a kakekotoba (pivot word) with a secondary meaning related to the sea?
Student: 塩/潮
Toki: Spot on. “しほ” functions as a pivot word meaning both “salt” (or the specific proper noun Shio) and “tide”, which connects perfectly with the associated words “磯” (rocky beach) and “千鳥” (plovers). Having established the structure and rhetorical devices, let us decipher the grammar by working from the end of the poem backwards. Looking at the final verb “なく”, what conjugation form (katsuyōkei) must it be in due to the presence of the bound particle “ぞ” earlier in the line?
Student: 連体形
Toki: Excellent. The bound particle “ぞ” triggers the kakari-musubi rule, forcing the final verb “なく” into the rentaikei (attributive form) instead of the standard terminal form. Moving right to left, the next word is the particle “と” in the phrase “やちよと”. Based on its position immediately following “やちよ” (eight thousand generations) and modifying the verb “なく”, what is the grammatical function of this particle?
Student: Citational
Toki: Yes, it functions as a quotative particle indicating what the plovers are singing. Moving backwards to the preceding phrase “きみがみ世をば”, what is the grammatical function of the particle “ば” when attached to the case particle “を”?
Student: strong emphasis
Toki: Precisely. The addition of the binding particle “は” (voiced here as “ば” due to sequential voicing) to the accusative case particle “を” serves to heavily

emphasize the direct object. Moving further leftward, we encounter the phrase “きみがみ世”. Connecting the noun “きみ” (lord or emperor) with the noun “み世” (reign or age), what is the grammatical function of the particle “が” here?
Student: possessive
Toki: Correct. The particle “が” functions as a possessive or genitive case marker, linking “きみ” and “み世” to mean “our lord’s reign.” Continuing leftward, we encounter the phrase “すむ千鳥”. Since the verb “すむ” (to live or dwell) is directly modifying the noun “千鳥” (plovers), what conjugation form (katsuyōkei) must “すむ” be in?
Student: 連体形
Toki: Correct. Because it directly modifies the noun “千鳥”, “すむ” takes the rentaikei (attributive form). Moving to the preceding phrase “さしでのいそに”, what is the grammatical function of the particle “に” attached to the noun “いそ” (rocky beach)?
Student: Location
Toki: Correct. The case particle “に” specifies the spatial location where the action of dwelling occurs. Moving further left to the phrase “さしでのいそ”, what is the grammatical function of the particle “の” linking the proper noun “さしで” to “いそ”?
Student: Possessive
Toki: Correct. The particle “の” acts as a genitive or possessive marker, linking the location “さしで” to the “いそ” (rocky beach). Moving to the very first phrase of the poem, “しほの山”, what is the grammatical function of the particle “の” connecting “しほ” to “山”?
Student: Same as before
Toki: Correct. Just as before, the particle “の” functions as a genitive marker, linking the name “しほ” to “山”. We have now successfully dissected the grammar and structure of the entire poem from end to beginning. Taking into account all the elements we have analyzed—the plovers crying “eight thousand generations” at the rocky beach of Sashide by Mount Shio, the strong emphasis on the lord’s reign, and the pivot word connecting the salt/tide to the coastal scene—how would you paraphrase the overall meaning of this waka in your own words?
Student: By Mt. Shio / on the Sashide Shore / dwell the plovers / “May our lord’s reign last eight-thousand years,” / they cry
Toki: That is an excellent paraphrase. You have perfectly captured the grammatical structures, the locational setting, and the strong emphasis on the lord’s reign. Could you please provide the next section of our text? I get tired after a while, so please open a new window for this new section.

Log 2. Toki (Beginner Student) [*Kokinshū* 21]
Student: 君がためはるののにいでてわかなつむわが衣手に雪はふりつつ
Toki: This text is a waka poem. How would you segment the syllables of this text into the traditional 5-7-5-7-7 structure?
Student: 君がためは/るののにいでて/わかなつむ/わが衣手に/雪はふりつつ
Toki: Take a closer look at your first and second segments. The characters "はる" form the word for "spring" (春). Keeping this in mind, how would you adjust the boundary between the first and second segments?
Student: 君がため/はるののにいでて/わかなつむ/わが衣手に/雪はふりつつ
Toki: Spot on. You correctly identified "はる" (春) and adjusted the boundary. You might notice the second segment (はるののにいでて) has eight syllables instead of the standard seven—this is a classic example of ji-amari (hypermetry) and your division is absolutely correct. Now that the structure is clear, let's begin our grammatical analysis by looking at the very end of the poem. The main predicate ends with "ふりつつ" (furi-tsutsu). Looking at the suffix "つつ" at the end of the sentence, what grammatical function and meaning does this conjunctive particle convey?
Student: That it is still going on? I forget the technical term.
Toki: You have nailed the exact meaning; it indicates that the action is ongoing or continuing. Since this particle generally serves to link clauses together (even though it is left trailing at the end of this specific poem for a lingering poetic effect), what specific category of particle (joshi) does "つつ" belong to?
Student: The one that links two phrases
Toki: Exactly. The technical term for a particle that connects clauses like this is a conjunctive particle (接続助詞, setsuzokujoshi). Now that we have established the function of "つつ", let's move one step to the left and look at the word immediately preceding it: "ふり" (furi), which comes from the verb "ふる" (furu, to fall). Knowing that the conjunctive particle "つつ" always dictates the form of the verb it attaches to, what conjugation form (活用形, katsuyōkei) must "ふり" be in this sentence?
Student: mizenkei?
Toki: Close! Let's look closely at the final vowel of "ふり" (fu-ri). If this were the Mizenkei (imperfective form) of this verb, it would end in an "a" sound, making it "ふら" (fu-ra). Since "ふり" ends in an "i" sound and is used here to connect to the conjunctive particle "つつ", what conjugation form must it be?
Student: renyōkei
Toki: Excellent. It is indeed the Renyōkei (continuative form). Moving one step further to the left, we encounter "雪は" (yuki wa). "雪" is a noun meaning "snow." What grammatical role does the particle "は" (wa) play in relation to this noun
Student: subject

Toki: Exactly. It marks “雪” (snow) as the subject of the falling action. Moving left, we encounter the phrase “衣手に” (koromode ni). Based on the kanji “衣” (garment) and “手” (hand), what specific part of the clothing does the poetic noun “衣手” refer to?
Student: The sleeves??
Toki: Spot on. “衣手” (koromode) elegantly refers to the sleeves of a robe. Now, let’s look at the particle “に” (ni) attached to it. What grammatical function does this particle serve in relation to the falling snow?
Student: “On” as in: the snow is falling (and accumulating) on the sleeves
Toki: Spot on. The particle “に” indicates the target or location of the action—the snow is falling onto the sleeves. Moving one step further to the left, we find the phrase “わが” (waga) directly preceding “衣手”. “わ” is a first-person pronoun meaning “I”. In Classical Japanese, what grammatical function does the particle “が” (ga) serve when it connects a pronoun like “わ” to a noun like “衣手”?
Student: I think possessive
Toki: Exactly right. It functions as a possessive marker, making “わが衣手” mean “my sleeves.” Putting this back half of the poem together, we have the image of snow continuously falling on the speaker’s sleeves. Moving one step further to the left, we find the phrase “わかなつむ” (wakana tsumu). The word “つむ” (tsumu) is a verb meaning “to pluck” or “to pick.” Notice that “つむ” directly precedes and modifies the noun phrase “わが衣手” (specifically modifying the “わ” or “I” whose sleeves they are). In Classical Japanese, when a verb directly modifies a noun that comes immediately after it, what conjugation form (katsuyōkei) must that verb be in?
Student: Can you remind me of my options?
Toki: Here are the six classical conjugation forms (活用形):
Mizenkei (Imperfective / Irrealis)
Renyōkei (Continuative)
Shūshikei (Terminal / Conclusive)
Rentaikei (Attributive)
Izenkei (Perfective / Realis)
Meireikei (Imperative)
Out of these six, which form is specifically used when a verb modifies a noun (substantive) that comes immediately after it?
Student: Rentaikei
Toki: Spot on. It is indeed the Rentaikei (attributive form), which literally means the “form that connects to a substantive (noun).” Here, the phrase “わかなつむ” modifies the speaker (“わ”) whose sleeves are gathering snow. Moving one step to the left, we find the noun “わかな” (wakana), which translates to “young greens” or “spring herbs.” Given the action of the verb “つむ” (to pick)

immediately following it, what grammatical role does “わかな” play in relation to this verb, even though a specific particle is omitted here for poetic rhythm?
Student: Is the particle that is omitted を?
Toki: Spot on. The implied particle is indeed “を” (wo), which marks “わかな” (young greens) as the direct object of the verb “つむ” (to pick). Moving one step further to the left, we encounter the word “いでて” (idete). This word is composed of a verb (“いで”) and the conjunctive particle “て” (te). Focusing on that particle “て”, what grammatical function does it generally serve when linking clauses or actions together?
Student: It indicates that the action continues
Toki: You might be thinking of “つつ” from a moment ago! While both are indeed conjunctive particles used to link clauses, “て” serves a slightly different primary function. Think about how the -te form is used to connect verbs even in modern Japanese (for example, “I woke up and ate breakfast”). How would you describe the temporal relationship or sequence that “て” creates between the action that precedes it and the action that follows it?
Student: “and”
Toki: Exactly. It functions just like “and” or “and then,” indicating a simple sequence of actions—first doing one thing, and then doing the next. Now let’s look at the verb “いで” (ide) attached to it. The dictionary form of this verb is “いづ” (izu), meaning “to go out” or “to come forth.” Just like the particle “つつ” we looked at earlier, the conjunctive particle “て” always requires the verb preceding it to be in a specific conjugation form. Based on that rule, what form (katsuyōkei) must “いで” be in?
Student: I think it is the continuative renyōkei
Toki: Spot on. It is indeed the Renyōkei (continuative form). Moving one step further to the left, we find the particle “に” (ni) attached to the noun phrase “はるのの” (haru no no, meaning “the fields of spring”). Thinking about the verb “いで” (to go out) that immediately follows it, what grammatical function does this particle “に” serve here?
Student: Location or direction of action
Toki: Spot on. It marks the destination—going out to or into the fields of spring. Moving one step further to the left, we look at the phrase “はるのの” (haru no no). As you noted when we first segmented the poem, “はる” means “spring.” The second “の” represents the noun “野” (field). What is the grammatical function of the first “の” that connects “はる” and “野”?
Student: Possessive
Toki: Exactly. It acts as a possessive or modifying marker, giving us “the fields of spring.” Moving to the very beginning of the poem, we have the phrase “君がため” (kimi ga tame). The word “君” (kimi) means “you,” and “ため” (tame)

means “sake” or “benefit.” Based on our earlier analysis of the phrase “わが衣手” (waga koromode), what grammatical function does the particle “が” (ga) serve here between “君” and “ため”?
Student: Also possessive
Toki: Spot on. It functions exactly like the “が” in “わが”, acting as a possessive to mean “for your sake.” Now that we have successfully worked our way backward through every part of the poem and deciphered all the grammatical components, how would you paraphrase the overall meaning of this poem in your own words?
Student: For you, I go on the spring fields to pick young herbs, but snow keeps falling on my robes
Toki: That is an excellent and highly accurate paraphrase. You perfectly captured the continuous action of the snow and the precise grammatical connections we uncovered throughout the poem. What is the next section of text you would like to analyze? I get tired after a while, so please open a new window for this new section.

Log 3. Toki (*kanbun*) [*Han Feizi*]
Student: 楚人有鬻盾與矛者
Toki: Looking at the original Chinese syntactic structure (Subject-Verb-Object), which character functions as the main predicate indicating existence in this sentence?
Student: 有
Toki: That is correct. In the original Chinese word order, what is the entire noun phrase that serves as the grammatical object of “有” (the entity that exists)?
Student: 鬻盾與矛者
Toki: That is correct. Within the phrase “鬻盾與矛者”, which character functions as the verb?
Student: 鬻
Toki: That is correct. Within this phrase, which characters serve as the direct objects of the verb “鬻”?
Student: I am not sure
Toki: In classical Chinese syntax, the grammatical object typically follows the verb directly. Which character immediately follows the verb “鬻” to indicate the first item being sold?
Student: 盾
Toki: That is correct. Which character represents the second item being sold in this phrase?
Student: 與
Toki: The character “與” functions as a coordinating conjunction connecting the two nouns in this list. Which character immediately follows this conjunction to

represent the second item being sold?
Student: 矛
Toki: That is correct. Following Japanese word order, how would you rearrange the characters within the phrase “鬻盾與矛者”?
Student: 盾と矛とを鬻ぐ者（I looked up 鬻 in the dictionary and got ひさぐ）
Toki: That is correct. Taking the subject “楚人”, the main verb “有”, and the entire object phrase you just assembled (“盾と矛とを鬻ぐ者”), how would you arrange these three main components into standard Japanese word order?
Student: 楚人（の中）に盾と矛とを鬻ぐ者あり。
Toki: That is correct. You have successfully transcribed the text into its Japanese reading form (kakikudashibun). How would you paraphrase the overall meaning of this sentence in your own words?
Student: Among the people of So there are persons who trade in shields and spears
Toki: That is an accurate paraphrase. What is the next section of the text? I get tired after a while, so please open a new window for this new section.

Log 4. Toki [*Kinyōshū* 270]
Student: あはぢしまかよふちどりのなくこゑにいくよねざめぬすまのせきもり
Toki: Could you please segment this text into the traditional 5-7-5-7-7 structure of a waka?
Student: あはぢしま/かよふちどりの/なくこゑに/いくよねざめぬ/すまのせきもり
Toki: That segmentation is correct. Does this poem reference or subvert a specific classical poem or tale?
Student: I think The Tale of Genji, the Suma episode
Toki: That is correct. The poem draws heavily on the atmosphere of the Suma chapter in The Tale of Genji, specifically the loneliness of listening to the plovers. Following our process for waka, we will now proceed through the text from right to left. Looking at the very end of the poem, what does the noun “sekimori” (せきもり) mean?
Student: It is some kind of guard at a mountain pass. But I don’t get what this has to do with the shore.
Toki: You are absolutely right that “sekimori” means a barrier guard. While we often picture mountain passes, checkpoints or barriers (seki) were also placed along important coastal routes, and the barrier at Suma was historically significant. Moving one step to the left, what is the grammatical function of the particle “no” (の) connecting the noun “Suma” (すま) to “sekimori” (せきもり)?
Student: Possessive? Or Locative? This is the guard at Suma
Toki: You have the right idea. It acts as an attributive/genitive marker linking the two nouns, practically indicating the guard located at Suma. Moving left to the word immediately preceding “Suma”, we encounter the suffix “nu” (ぬ)

attached to the verb “nezame” (ねざめ). Is this “nu” acting as the perfective auxiliary (完了) or the negative auxiliary (打消)?
Student: If it is negative it should follow mizenkei, and if it is perfective instead renyōkei, right? But nezamu is a shimo-nidan verb, so both conjugations are nezame. How am I supposed to know which is the case?
Toki: That is an excellent observation. You are exactly right that for a shimo-nidan verb like nezamu, both the mizenkei and ren’yōkei are identical, which creates a deliberate structural ambiguity here. To resolve this, we must look at how “nu” relates to the rest of the poem. If “nu” were directly modifying the noun phrase “Suma no sekimori” that follows it, it would need to be in the attributive form (rentaikei). What are the rentaikei forms for both the negative auxiliary “zu” and the perfective auxiliary “nu”?
Student: Right! If it were perfective the rentaikei would be nuru. It has to be negative.
Toki: Spot on. Because “nu” directly modifies the noun phrase “Suma no sekimori,” it must be in the rentaikei form. As you correctly deduced, the rentaikei of the negative auxiliary “zu” is “nu,” confirming that it is indeed negative. Moving further to the left, we find the phrase “iku yo” (いくよ). What does this phrase mean?
Student: how many nights? how many generations?
Toki: You have correctly identified the two possibilities! “Iku yo” can indeed be written as 幾夜 (“how many nights”) or 幾代 (“how many generations/ages”). Considering the adjacent verb “nezame” (waking from sleep), which of these two meanings fits the immediate context better?
Student: Nights, probably.
Toki: Exactly. “How many nights” makes perfect sense in conjunction with the act of waking up. Continuing leftwards, we encounter the phrase “naku koe ni” (なくこゑに). What is the grammatical function or meaning of the case particle “ni” (に) here as it connects to the later action of waking (“nezame”)?
Student: It marks the event that woke the guard up
Toki: That is correct. It functions as a particle of cause or reason, indicating why the guard wakes up. Moving left, we find the phrase “chidori no” (ちどりの). What is the grammatical function of the particle “no” (の) here as it relates to the verb “naku” (なく)?
Student: marks the subject of the verb
Toki: Exactly. In Classical Japanese, the particle “no” often functions like the modern “ga” to mark the subject of a clause—in this case, the plovers doing the crying. Moving left to the verb “kayofu” (かよふ), what conjugation form must it be in since it directly precedes the noun “chidori” (ちどり)?
Student: Rentaikei
Toki: Spot on. Because it directly modifies the noun “chidori,” it is in the rentaikei (attributive form). Finally, we reach the very beginning of the poem:

"Awajishima" (あはぢしま). Based on the meaning of the verb "kayofu" (to travel back and forth / to frequent), what is the relationship between Awaji Island and the plovers?
Student: The plovers come and go?
Toki: Yes, exactly. They fly back and forth to Awaji Island. Now that we have successfully analyzed every component of the poem, could you put all these pieces together and paraphrase the overall meaning of the poem in your own words?
Student: Back and forth fly / the plovers to Awaji Island / and their voices / how many night have they awakened / the guard at Suma?
Toki: Your paraphrase captures the poetic atmosphere wonderfully. However, look closely at how you translated the action of the guard: "have they awakened." This phrasing treats the verb as a completed, positive action. Based on our established conclusion that the auxiliary "nu" (ぬ) here is the negative form, how should we adjust the English paraphrase to reflect this negative meaning regarding his sleep?
Student: "haven't they awakened"
Toki: That translation perfectly matches the negative grammar we established. However, let's consider the context of the Suma chapter you identified earlier: if the guard "does not wake up," could he be listening to the sad cries of the plovers in the night?
Student: I am a bit lost here. How would you translate it?
Toki: You have identified the exact contradiction! Since the guard must be awake to hear the plovers, our previous assumption that "nu" is negative—and therefore modifying the following noun—must be incorrect. If the phrase "iku yo nezamenu" is actually a complete statement ending right there, what conjugation form must "nu" take?
Student: I see! How about this? "Back and forth fly / the plovers to Awaji Island / and their voices / how many night have they not awakened him? / The guard at Suma"
Toki: Your logic for incorporating the negative is structurally sound based on our earlier assumption! However, as we noticed, a sleeping guard cannot hear the sad cries of the plovers. This tells us our structural assumption—that "nu" directly modifies "Suma no sekimori"—was a trap. In waka poetry, it is very common to use inversion (tōchihō), where the grammatical sentence ends early, and the subject or a related noun is tacked on at the very end for emotional emphasis. If the grammatical sentence actually ends at "nezamenu", then "nu" must be in the terminal form (shūshikei). Between our two candidates—the negative auxiliary "zu" and the perfective auxiliary "nu"—which one has "nu" as its shūshikei?
Student: Perfective
Toki: Correct. Because "nu" is the shūshikei of the perfective auxiliary, the grammatical sentence ends there, indicating the completed action of waking up.

The phrase “Suma no sekimori” is an inverted subject appended at the end for emotional resonance.
Toki: Now that we have successfully deciphered all the grammar and resolved this structural inversion, how would you paraphrase the final, accurate meaning of the entire poem in your own words?
Student: Back and forth fly / the plovers to Awaji Island / and their voices / countless nights they woke up / the guard at Suma?
Toki: That is an excellent paraphrase. You have successfully captured the perfective aspect of the action and the emotional weight of the inverted subject at the end of the poem. Please provide the next section. I get tired after a while, so please open a new window for this new section.

Log 5. Hiki (Classical Japanese to English)
Student: 行く
Hiki: ゆく 【行く・往く】
〔動詞・カ行四段活用〕（ゆか・ゆき・ゆく・ゆく・ゆけ・ゆけ）To move from one place toward another; to proceed, depart, or pass on.
〔自動詞〕 Indicates physical or temporal movement away from the speaker’s standpoint.
(1) To go; to travel toward a destination.「東の方に住むべき国求めにとて行きけり」（伊勢物語・九段）
(2) To depart; to leave a place behind.「君が行く道の長手を繰り畳ね焼き滅ぼさむ天の火もがも」（万葉集・三七二四）
(3) (Of time, seasons, or years) to pass; to elapse.「行く春を近江の人と惜しみける」（猿蓑）
(4) (Of water, clouds, etc.) to flow on; to drift along.「ゆく河の流れは絶えずして、しかももとの水にあらず」（方丈記）
(5) To die; to pass away (euphemistic).
(6) (As 「心ゆく」) to be satisfied; to feel at ease; for one’s feelings to be soothed.

〔補助動詞〕 Attached to the 連用形 of a verb (often via 「て」), indicating that an action continues onward or that a change progresses gradually.
(1) To go on doing; to become increasingly so.「日数の早く過ぐるほどぞ、ものにも似ぬ。…おぼつかなさもまさり行く」
〔参考〕「ゆく」と「いく」は古くから併用されるが、和歌・雅文では「ゆく」が正格とされ、「いく」はやや口語的・俗語的な形とされる。

Student: Would any of these expressions work also for the meaning “pilgrimage”?

Hiki: Not really — 行く alone simply denotes going, so for pilgrimage premodern Japanese prefers 詣づ（まうづ） or 参る, humble verbs specifically used for visiting shrines and temples. One could say 物詣でに行く, but the pilgrimage sense there comes from 物詣で, not from 行く itself.
Student: Is there a distinction between a pilgrimage to a temple versus to a shrine?
Hiki: Not in the verb itself — 詣づ and 参る serve for both Buddhist temples and Shinto shrines, reflecting the premodern blending of the two traditions (神仏習合).

Log 6. Hiki (English to Classical Japanese)
Student: pilgrimage
Hiki: ものまうで 【物詣で】〖名〗 A pilgrimage; the act of going to pray or worship at a temple or shrine (e.g., Hatsuse, Ishiyama, Kiyomizu). Frequently performed by Heian-period courtiers and ladies. 「ものまうでし給ひけるに」(源氏物語)
まうで 【詣で】〖名〗 A visit of worship to a shrine, temple, or to a person of high rank; pilgrimage. (Nominal form of the verb まうづ.)
じゅんれい 【巡礼・順礼】〖名・サ変〗 A pilgrimage made by traveling in sequence to a series of sacred sites (such as the thirty-three Kannon temples of the western provinces); also, the pilgrim who undertakes such a journey.
Student: Is there a specific expression for a pilgrimage to Ise?
Hiki: Yes — 伊勢参り (いせまいり) or 伊勢参宮 (いせさんぐう), referring specifically to a worship visit to the Ise Grand Shrine; the older verbal expression is 伊勢に詣づ. The mass popular pilgrimages of the Edo period were also called おかげ参り.

Log 7. A Conversation with Sata
User: いづくへおはしますぞ。
Sata: 上の御局へまゐらむとて。いと急ぎの御使ひにさぶらふなり。
User: しばし、いとゆかしき事のありげなるかな、すこし語り給へ。
Sata: なべての事にはあらずかし。帰り参りて後にこそ、忍びて聞こえさせめ。
User: せめて、その事の端ばかりだに語り給へ。
Sata: あな、わりなや。さるやんごとなき方の御文なりとばかりにて、今は許し給へ。
User: もし、入道の后より院へ聞こえ給ふ御文にや。
Sata: あな、おそろし。さやうの事、ゆめゆめ人にな聞かせ奉り給ひそ。今はとて、いそぎ参り侍らむ。
User: Where are you heading?
Sata: I am going to the upper chambers. I am on a very urgent errand.

User: Wait a moment. It seems there is something fascinating going on. Please, do tell.
Sata: This is no ordinary matter. After I return, I shall tell you in confidence.
User: Please give me at least a hint.
Sata: Oh, how unreasonable you are! A letter, from someone in a very high place. Now, if you will excuse me.
User: Could it be a letter from the Nun Empress to the Cloistered Emperor?
Sata: Oh, how terrifying! About this matter, you should not say anything to anybody at all. And now, I must hurry on my way.

NOTES

Author Contributions: (1) Conceptualization, Methodology, Writing (original draft, review and editing). (2) Validation (linguistics and pedagogy), Writing (review and editing). (3) Methodology (prompt engineering), Writing (review and editing). (4) Investigation (user testing and feedback). Thanks also to Momoyo Kubo Lowdermilk, Ekaterina Mozhaeva, Małgorzata Citko-DuPlantis, John Groschwitz, Carolyn Zou, Ron Egan, John Mitchell, Alfonso Amat, Thomas McAuley, Hidenori Jinno, and Shigeki Moro.

REFERENCES

Aikawa, Takako. 2026. “AI-Enhanced Japanese Language Class with Customized Large Language Models.” Conference Presentation. In *SIG Panel: Artificial Intelligence*. AATJ 2026, Vancouver.

Alansari, Aisha, and Hamzah Luqman. 2026. “Large Language Models Hallucination: A Comprehensive Survey.” *arXiv*. https://arxiv.org/abs/2510.06265v3. Accessed September 19, 2026.

Borges, Jorge Luis. 1942. “Funes el memorioso.” *La Nación* (Buenos Aires), 7 June 1942, “Artes y Letras.”

———. 1964. *Labyrinths: Selected Stories & Other Writings*. Edited by Donald A. Yates and James E. Irby. New York: New Directions.

———. 1998. *Collected Fictions*. Translated by Andrew Hurley. New York: Viking.

Boudourides, Moses. 2026. “Structural Hallucination in Large Language Models: A Network-Based Evaluation of Knowledge Organization and Citation Integrity.” *arXiv*. https://arxiv.org/abs/2603.01341v1. Accessed September 19, 2026.

Burge, Marjorie, Jeffrey Niedermaier, and Pier Carlo Tommasi. 2025. “Contesting the ‘Classical,’ Creating Communities: The ‘Intercollegiate Classical Japanese

Poetry Contest' Within the Landscape of 2020s North American Bungo Pedagogy." *Japanese Language and Literature* 59 (2): 307–49. https://doi.org/10.5195/jll.2025.373. Accessed September 19, 2026.

Chau, Herman, Michael R. Zeng, and Paul S. Atkins. 2025. "Deriving Orthographic Data from Classical Japanese Texts with Machine-Learning Methods." In *JINMONKON 2025: Proceedings of the Annual Conference of the Computers and Humanities Special Interest Group*, Information Processing Society of Japan. https://hdl.handle.net/1773/55355. Accessed September 19, 2026.

Crawcour, Sydney. 1965. *An Introduction to Kambun*. Ann Arbor, M.I.: Center for Japanese Studies, University of Michigan. https://archive.org/details/crawcour-introduction-to-kambun. Accessed September 19, 2026.

Danyaro, Kamaluddeen Usman, Shamsu Abdullahi, Abdallah Saleh Abdallah, and Haruna Chiroma. 2025. "Hallucinations in Large Language Models for Education: Challenges and Mitigation." *International Journal of Teaching, Learning and Education* 4 (6): 13–19. https://doi.org/10.22161/ijtle.4.6.2. Accessed September 19, 2026.

Gabrovšek, Rok, and David Rihtaršič. 2025. "Custom Generative Artificial Intelligence Tutors in Action: An Experimental Evaluation of Prompt Strategies in STEM Education." *Sustainability* 17 (21). https://www.mdpi.com/2071-1050/17/21/9508. Accessed September 19, 2026.

Gladstone, Clovis, Zhao Fang, and Spencer Dean Stewart. 2025. "Ground Truth Generation for Multilingual Historical NLP Using LLMs." *arXiv*. https://arxiv.org/abs/2511.14688v1. Accessed September 19, 2026.

Haddow, Barry, Rachel Bawden, Antonio Valerio Miceli Barone, Jindřich Helcl, and Alexandra Birch. 2022. "Survey of Low-Resource Machine Translation." *Computational Linguistics* 48 (3): 673–732. https://direct.mit.edu/coli/article/48/3/673/111479/Survey-of-Low-Resource-Machine-Translation. Accessed September 19, 2026.

Itō, Rika. 2026. "Generating Student Agency: A Report on Integrating Generative AI into Third-Year Japanese Writing for Metacognitive and Language Development." Conference Presentation. In *Language and Technology Papers*. AATJ 2026, Vancouver.

Iwabuchi, Etsutarō. 1972. *Bungo no bunpō*. Tokyo: Shūei shuppan.

Komai, Akira, and Thomas Rohlich. 1988. *An Introduction to Japanese Kanbun*. Nagoya: University of Nagoya Press.

Kumar, Harsh, Zi Kang (Jace) Mu, Jonathan Vincentius, and Ashton Anderson. 2026. “Beyond the AI Tutor: Social Learning with LLM Agents.” *arXiv*. https://arxiv.org/abs/2604.02677. Accessed September 19, 2026.

Laban, Philippe, Hiroaki Hayashi, Yingbo Zhou, and Jennifer Neville. 2026. “LLMs Get Lost in Multi-Turn Conversation.” In *Proceedings of the Fourteenth International Conference on Learning Representations*. https://arxiv.org/abs/2505.06120. Accessed September 19, 2026.

Lege, Ryan, Euan Bonner, and Takako Aikawa. 2024. “Pecha, a Language Practice Peer: Guiding Language Learning Interactions Through Large Language Models.” *Technology in Language Teaching and Learning* 6 (3): 1–14. https://www.castledown.com/journals/tltl/article/view/tltl.v6n3.1716. Accessed September 19, 2026.

Li, Haolin, Haipeng Zhang, Mang Li, et al. 2026. “Toward Robust Multilingual Adaptation of LLMs for Low-Resource Languages.” *arXiv*. https://arxiv.org/abs/2510.14466. Accessed September 19, 2026.

Liu, Jiayu, Zhenya Huang, Tong Xiao, et al. 2024. “SocraticLM: Exploring Socratic Personalized Teaching with Large Language Models.” *Proceedings of the 38th Conference on Neural Information Processing Systems (NeurIPS)*. https://proceedings.neurips.cc/paper_files/paper/2024/file/9bae399d1f34b8650351c1bd3692aeae-Paper-Conference.pdf. Accessed September 19, 2026.

McCullough, Helen Craig. 1988. *Bungo Manual: Selected Reference Materials for Students of Classical Japanese*. Ithaca, N.Y.: East Asia Program, Cornell University.

McPherson, Naemi. 2026. “Reframing Learning Through Translation: Fostering Critical Thinking and Awareness of Diversity in Advanced Japanese Education in the Age of AI.” Conference Presentation. In *Language and Technology Papers*. AATJ 2026, Vancouver.

Nichieisha Editorial Department. 1998. *Shin kanbun no kihon nōto*. Tokyo: Nichieisha.

Nijōan Shujin [Nobuyuki Kaji]. 2010. *Kanbunpō kiso: Hontō ni wakaru kanbun nyūmon*. Tokyo: Kōdansha.

Ōno, Susumu, et al. 2006. *Shinpen bungo bunpō*. Kyoto: Chūō tosho shinsha.

Özateş, Şaziye Betül, Tarık Emre Tıraş, Ece Elif Adak, et al. 2025. “Building Foundations for Natural Language Processing of Historical Turkish: Resources and Models.” *arXiv*. https://arxiv.org/abs/2501.04828. Accessed September 19, 2026.

Shirane, Haruo. 2005. *Classical Japanese: A Grammar*. New York: Columbia University Press.

———. 2007. *Classical Japanese Reader and Essential Dictionary*. New York: Columbia University Press.

Shu, Peng, Junhao Chen, Zhengliang Liu, et al. 2024. "Transcending Language Boundaries: Harnessing LLMs for Low-Resource Language Translation." *arXiv*. https://arxiv.org/abs/2411.11295. Accessed September 19, 2026.

Takahashi, Jun. 2026. "How Do Beginner Learners of Japanese Use ChatGPT in the Writing Process: An Engagement Perspective." Conference Presentation. In *Language and Technology Papers*. AATJ 2026, Vancouver.

Takeuchi, Jae DiBello. 2026. "AI Voice Mode for Intermediate JFL Learners: Pros, Cons, and Pragmatic Pitfalls." Conference Presentation. In *SIG Panel: Artificial Intelligence*. AATJ 2026, Vancouver.

Tomono, Takao, and Yumi Sugihara. 2026. "Beyond the Dual Allure and Dual Dispossession of Generative AI and Native-Speaker Norms: An Examination of a Japanese Academic Writing Class That Cultivates Learner Agency and Voice." Conference Presentation. In *Language and Technology Papers*. AATJ 2026, Vancouver.

Tsuchiya, Shinsuke, and Brigham Graham. 2026. "Turning AI into a Helpful Coach: Prompt-Giving and Codeswitching for Targeted Speaking Practice." Conference Presentation. In *SIG Panel: Artificial Intelligence*. AATJ 2026, Vancouver.

Wang, Hao, Hirofumi Shimizu, and Daisuke Kawahara. 2023. "Kanbun-LM: Reading and Translating Classical Chinese in Japanese Methods by Language Models." In *Findings of the Association for Computational Linguistics: ACL 2023*, 8589–8601. Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.findings-acl.545. Accessed September 19, 2026.

Wixted, John Timothy. 2006. *A Handbook to Classical Japanese*. Ithaca, N.Y.: East Asia Program, Cornell University.

Wood, David, Jerome S. Bruner, and Gail Ross. 1976. "The Role of Tutoring in Problem Solving." *Journal of Child Psychology and Psychiatry* 17: 89–100. https://doi.org/10.1111/j.1469-7610.1976.tb00381.x. Accessed September 19, 2026.

Zhong, Tianyang, Zhenyuan Yang, Zhengliang Liu, et al. 2024. "Opportunities and Challenges of Large Language Models for Low-Resource Languages in Humanities Research." *arXiv*. https://arxiv.org/abs/2412.04497. Accessed September 19, 2026.