Not One Prompt

Making a modern-Bengali edition of Swami Saradananda's *Śrīśrīrāmakṛṣṇa-Līlāprasaṅga* with reviewed AI agents

A project report

Tamal Maharaj

Department of Computer Science

Ramakrishna Mission Vivekananda Educational and Research Institute, Belur Math

tamal@gm.rkmvu.ac.in

November 2026

Abstract

Swami Saradananda's *Śrīśrīrāmakṛṣṇa-Līlāprasaṅga* is one of the most searching accounts of Sri Ramakrishna's life and spiritual practice, but it is written in *sādhu bhāṣā*, the formal literary Bengali of the early twentieth century, which many readers today find hard going. We made a version of the whole book in modern standard Bengali (*calit*) that keeps the author's own words wherever they still work. Every sentence is linked to the original, so a reader can see the author's words with one click. The work was done by AI agents, but not by asking a chatbot to "translate" the book. Six agents with different jobs passed every one of the book's 16,838 sentences in turn. They were a translator, a meaning reviewer, a literary critic, a reviser, a deliberately sceptical reader and a final reviser. They worked from a style guide built from the editor's own rulings, a 2,037-word vocabulary, a list of 280 people and other named figures, and the published English translation. The editor ruled on the questions they raised, and machine checks guarded what must never change. This report describes the method and what each stage did, with numbers measured from the project's files and logs. It also gives the cost (700 agent runs, 92 agent-hours) and what the method still cannot guarantee. More than one sentence in four of a carefully prepared first draft was changed by review. The faults the reviewers caught are exactly the kind a one-line prompt produces and a casual reader would not notice.

1 Introduction

Swami Saradananda (1865–1927), a direct disciple of Sri Ramakrishna, wrote the *Līlāprasaṅga* between 1909 and 1919 in five parts. It is at once a biography, an account of spiritual practice from someone who knew it at first hand, and a work of Bengali prose of real gravity. It is written in *sādhu bhāṣā*, the formal written Bengali of its time. A reader brought up on modern Bengali meets verb forms, pronouns and long participial sentences that slow every line. For readers whose first language is not Bengali but who wish to read the book in Bengali, the barrier is higher still. Many begin the book and set it aside because of the language.

This project made a modern rendering of the whole book, the **Saral Līlāprasaṅga** (*saral*, "simple"), now published at

`https://cs.rkmvu.ac.in/~tamal/saral-lilaprasanga.html`

It is not a retelling and not a translation into another language. It is the same book in the register of today's written Bengali, with the author's vocabulary, names, scriptural quotations, songs and verses kept. The reader can click any sentence to see the original beneath it, or read the original and open the modern sentence only where needed.

The work was done with AI agents (Claude Opus 5.5, made by Anthropic) under the direction of an editor, the author of this report. Hence this report. Many people's picture of AI translation is a single instruction such as "convert this into *calit*". When the result disappoints, they conclude that the tool is poor. Our experience was

different. The instruction was the smallest part of the work. What made the result usable was everything around it:

- a precise definition of the target, drawn from the editor's own verdicts on sample sentences;
- knowledge prepared in advance: an English alignment, a list of people, a vocabulary and a style guide;
- several reviewers with different jobs and different blind spots, each of whom had to put any objection as a full corrected sentence;
- written reasons for every proposal that was rejected;
- an editor whose rulings reached every later stage;
- machine checks that guard what must not change, while the text itself is changed only by judgement, never by formula.

The rest of this report describes these parts, follows one sentence through every stage, and gives measured figures for what each stage did and what it cost. Section 12 sets out plainly what has not yet been done.

2 The task

2.1 Sādhu and calit

Written Bengali had two standard forms for most of the twentieth century. The *sādhu* form used full verb forms (*karitechilen*) and full pronouns (*tāṃhār*). The *calit* form was based on the speech of Calcutta (*karchilen*, *tāṃr*). Ferguson's classic account of diglossia [3] describes the pattern; Chatterji [4] traces its history. Today *calit* has replaced *sādhu* almost everywhere.

At first sight the conversion is a matter of word endings:

করিতেছিলেন → করছিলেন তাঁহার → তাঁর হইতে → থেকে ঐরূপ → এরকম

Several things make it much harder:

- *Sādhu* sentences carry their structure in grammar that *calit* has partly lost. Examples are case endings such as *-digake*, feminine agreement in adjectives (অভিজ্ঞা), and passive constructions. Convert these mechanically and the sentence can come to say something else (Section 3).
- Many words kept their spelling but changed their everyday sense. A modern reader hears today's sense.
- A sentence converted word by word is often grammatical and still not Bengali that anyone would write.
- The author's long periodic sentences, especially in the philosophical chapters, carry an argument. If its connectives are flattened, the argument breaks.

2.2 What "simple" means here

Before any text was produced, the editor read sample renderings and gave verdicts. These fixed the register, which we call *level A*: *calit* with nearly the same words. The verdicts showed two opposite ways to fail.

	Rejected	**Accepted**
Changed too much	কীর্তনের আনন্দ ("the joy of kīrtan")	কীর্তনানন্দ: the author's own rich word stays; a tap shows its meaning
Changed too little	সে বিষয়ে খোঁজ নিতে লেগেছিলেন	সে বিষয়ে খোঁজ করতে শুরু করলেন
Every word right, frame wrong	গদাধরের বয়সের অষ্টম বর্ষে ... সেভাবে অধিককাল ছিলেন	গদাধরের বয়স যখন আট বছর, তখন ... সেইরকমই অনেক দিন ধরে ছিলেন

Table 1: Three of the editor's early verdicts, which define the target. The third shows that a converted phrase can have every word defensible and still be something nobody would say: "nearly the same words" governs words, not the stiff shape of a phrase.

Register also follows the scene. Household matters (food, clothes, money, tears, boys' talk) go plain: the editor accepted খাবারদাবার and চোখের জল ফেলতে. Philosophy, devotion and samādhi keep the author's high style.

2.3 What must never change

Some things are fixed and every stage was told so. They are: verses, songs, mantras and scriptural quotations; footnote marks, numerals, dates and the sign ৺; names in the form the author used; and the level of respect (তাঁর for those the author honours). Qualifiers such as যেন ("as if"), প্রায় ("almost") and বোধ হয় ("perhaps") must stay, and so must every negation. The qualifiers matter more than they seem: যেন is how the author marks a visionary experience as reported rather than asserted. One sentence stays one sentence, and nothing is added or explained inside the text.

Sri Ramakrishna's own colloquial speech was at first to be copied exactly. After reading the result, the editor refined the rule: where printed speech mixes *sādhu* grammatical forms into colloquial sentences, those forms are converted and the author's spellings kept (116 sentences). This was one of several rulings made during the work (Section 7).

3 Why one prompt is not enough

The book has 16,838 sentences and about 280,000 words in 77 chapters. No single conversation with a model can hold it. Keeping choices consistent from the first chapter to the last, such as how a recurring word is rendered or who a recurring name refers to, needs a memory outside the model. In this project that memory was the style guide, the vocabulary, the list of people and a running ledger.

Scale is only half the problem. The other half is that the faults are *plausible*. Table 2 shows faults from the first draft. That draft was made by a translator agent that *had* the style guide, the vocabulary, the list of people and the English translation. Each fault reads as fluent Bengali. A reader of the modern text alone would not notice any of them, and that is the danger.

None of these is exotic. Each comes from something a mechanical conversion does naturally:

- it replaces an old case ending with its modern equivalent;
- it drops a feminine form that modern Bengali does not use;
- it turns a passive into an active;
- it keeps a word whose spelling survived but whose sense did not;
- it writes a gloss from the commonest sense of a word.

To catch these, someone has to read for meaning, sentence by sentence, against the original, and look up the rest of the book when in doubt. The meaning reviewer did exactly that. The same holds for style: the critic and the sceptical reader between them made over 3,800 findings of stiff, unclear or unidiomatic phrasing (Table 7).

4 Preparation

Four things were prepared before any sentence was converted (Figure 1).

The sentence table. The digital text of the book was split into 4,575 paragraphs and 16,838 sentences. Each has a stable id: `p-705.2` is the third sentence of paragraph 705. Most are body text; about 1,700 are headings and 423 are footnotes. Together they hold 279,702 words. Every later file refers to sentences by id. The book can be rebuilt from the table exactly, so nothing can be lost or duplicated without a check noticing.

The English alignment. Each Bengali sentence was paired with its counterpart in Swami Jagadananda's English translation, *Sri Ramakrishna the Great Master* [2], from the editor's earlier parallel reader of the two books. 14,419 Bengali sentences have an English counterpart. A further 565 English passages have no Bengali counterpart; these are mostly places where the translator expanded or rearranged the text. Every agent saw the English beside the Bengali, as a check on its reading, never as a source of words.

Kind of fault	What happened
A case ending flattened `p-1421.0`	Original: কামারপুকুরবাসী**দিগকে** ঠাকুরের অপূর্ব নূতনভাবে দেখিবার কারণ. "Why the Master saw the people of Kamarpukur in a wonderfully new way." Draft: কামারপুকুরবাসী**দের** ঠাকুরের ... দেখার কারণ. With two genitives, a reader today takes it the other way round: the villagers saw the Master anew. Caught by the meaning reviewer, graded *serious*. Final: কামারপুকুরবাসী**দেরকে** ...
Gender agreement lost `p-1148.1`	Original: ... সাধনে তিনিও বিশেষ **অভিজ্ঞা** ছিলেন. The feminine form tells the reader that "she" is the Brāhmaṇī. Modern Bengali drops the feminine adjective. The draft's অভিজ্ঞ, coming just after "from the Master's own lips", reads as if the Master himself was expert, which is wrong at this point in the story. Final: ... সাধনে **ব্রাহ্মণীও** বিশেষ অভিজ্ঞ ছিলেন.
Passive made active `p-705.2`	The original's passive (তাঁহার দ্বারা ... অনুরুদ্ধ হইয়া, "being requested by him") kept the subject clear. The draft's active এবং **তিনি** ... অনুরোধ করলে makes the Master appear to request himself. Followed in full in Section 6.
A word's sense has moved `p-519.2`	অজ্ঞাত দিব্যশক্তির **নিয়োগে**: in the original, "impelled by an unknown divine power". Today শক্তি নিয়োগ means *applying* power, so the sentence would say they themselves wield the power, which is the reverse of the meaning. Final: ... দিব্যশক্তির **প্রেরণায়**.
`p-664.2`	... মানবমনের ত্রুটিগুলিও তাঁকে যথাযথভাবে **স্বীকার** করতে হয়: here "take on" (English *assume*); today স্বীকার means "admit". God would be *confessing* the faults of the human mind rather than taking them on. Final: **গ্রহণ** করতে হয়.
A false fact in a gloss `p-3612.2`	The translator's tap-for-meaning note glossed বেটুয়া as "a small purse for money". The book itself says the Master's bag held spices (p-2726.5, p-3714.0), and the same chapter tells that his hand stiffened at the touch of metal. A reader tapping the word would learn something false. Caught by the meaning reviewer (*serious*): "a small bag for fennel, spices and the like".
Automatic glossing	An early attempt to underline hard words automatically from the vocabulary gave কারণ ("reason") the meaning "ritual wine" and চার ("four") the meaning "bait". Since then a meaning is shown automatically only for words with a single sense; every other note is written for its own sentence.
The English as a guide `p-3765.0`	*The Great Master* [2] helped at every step, but it compresses. Here it drops the author's point that other accomplished sādhakas had reached the same conclusion as the Master. Hence the rule: the Bengali decides, and the English only helps.

Table 2: Faults in the first draft and how they were caught. Sentence ids such as `p-1421.0` refer to paragraph 1421, sentence 0 of the book; the side-by-side reader opens at any id.

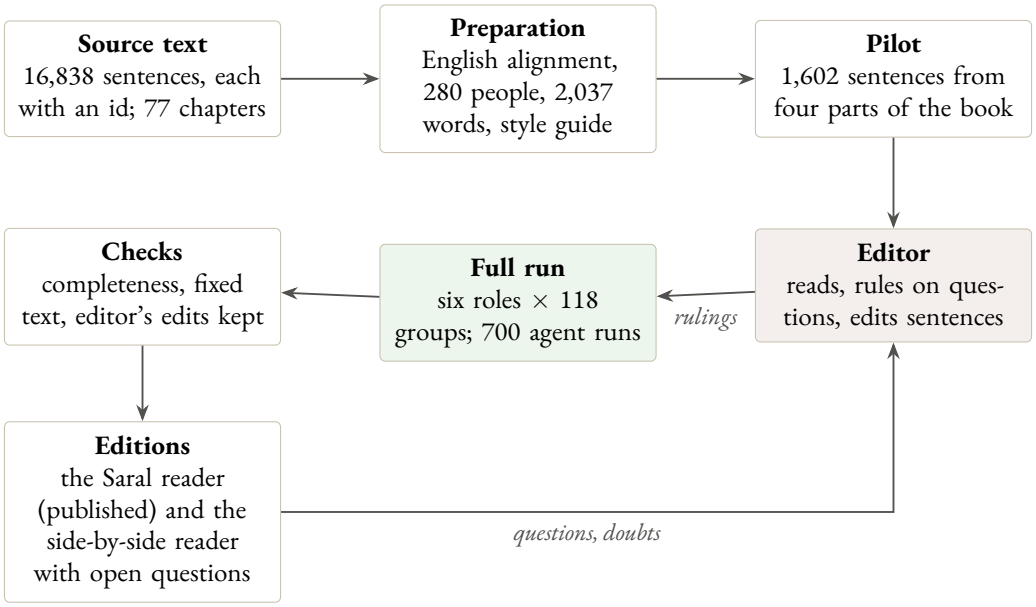


Figure 1: The project as a whole. The editor's rulings feed the style guide that every agent reads, and the questions the agents raise come back to the editor.

The people. The book has many people with the same names and many names for the same person. There are several Narens. মা may be the Holy Mother, Mother Kālī or someone's own mother. সারদা is only rarely the Holy Mother: elsewhere it is part of the name of a Tantra, or of Saradaprasanna. The list has 280 entries with 700 name forms. 82 of those forms are shared by more than one person, and 178 entries carry a warning. The notes printed at the top of each batch name the people who appear in it.

The vocabulary. 2,037 words were classified:

- 1,259 KEEP: a rich word stays, and the reader gets a tap-for-meaning note.
- 436 REPLACE: an obsolete form with its modern equivalent.
- 341 CONTEXT: decided sentence by sentence.

Some entries took real reading. ঠাকুরবাটী, for instance, occurs 60 times. All of them were read: it always means the whole temple establishment (shrines, courtyard, staff, stores, garden) and never the Master's house. One passage sets it against মা কালীর মন্দির ("Mother Kālī's temple"), so it can never be rendered মন্দির. Three spaced occurrences, ঠাকুর বাটী, are a different thing altogether: ঠাকুর (the Master, as subject) followed by বাটী ("house").

The style guide. The style guide is a 43-kilobyte document in eleven sections. It records the register, the fixed things, respect and voice, names, negation, sentence division, the use of the English, spelling, the output format, and two sections of lessons learned. It was written from the editor's verdicts and grew throughout the work. It is binding on every agent except one (Section 5).

Batches and groups. The book was cut into 153 batches. These were grouped into 118 groups of at most 225 sentences, each within a single chapter, so that every reviewer reads continuous text and can hear how a page sounds as a whole.

5 The pipeline

Every group passed six roles in order (Figure 2). Each role is a fresh AI agent with a short written brief of its own, read together with the shared references. Agents do not see one another's conversations; they communicate only through files. Each stage writes to its own folder and may not touch any other. A stage counts as finished only when a check program has verified its files; the files are then locked against change.

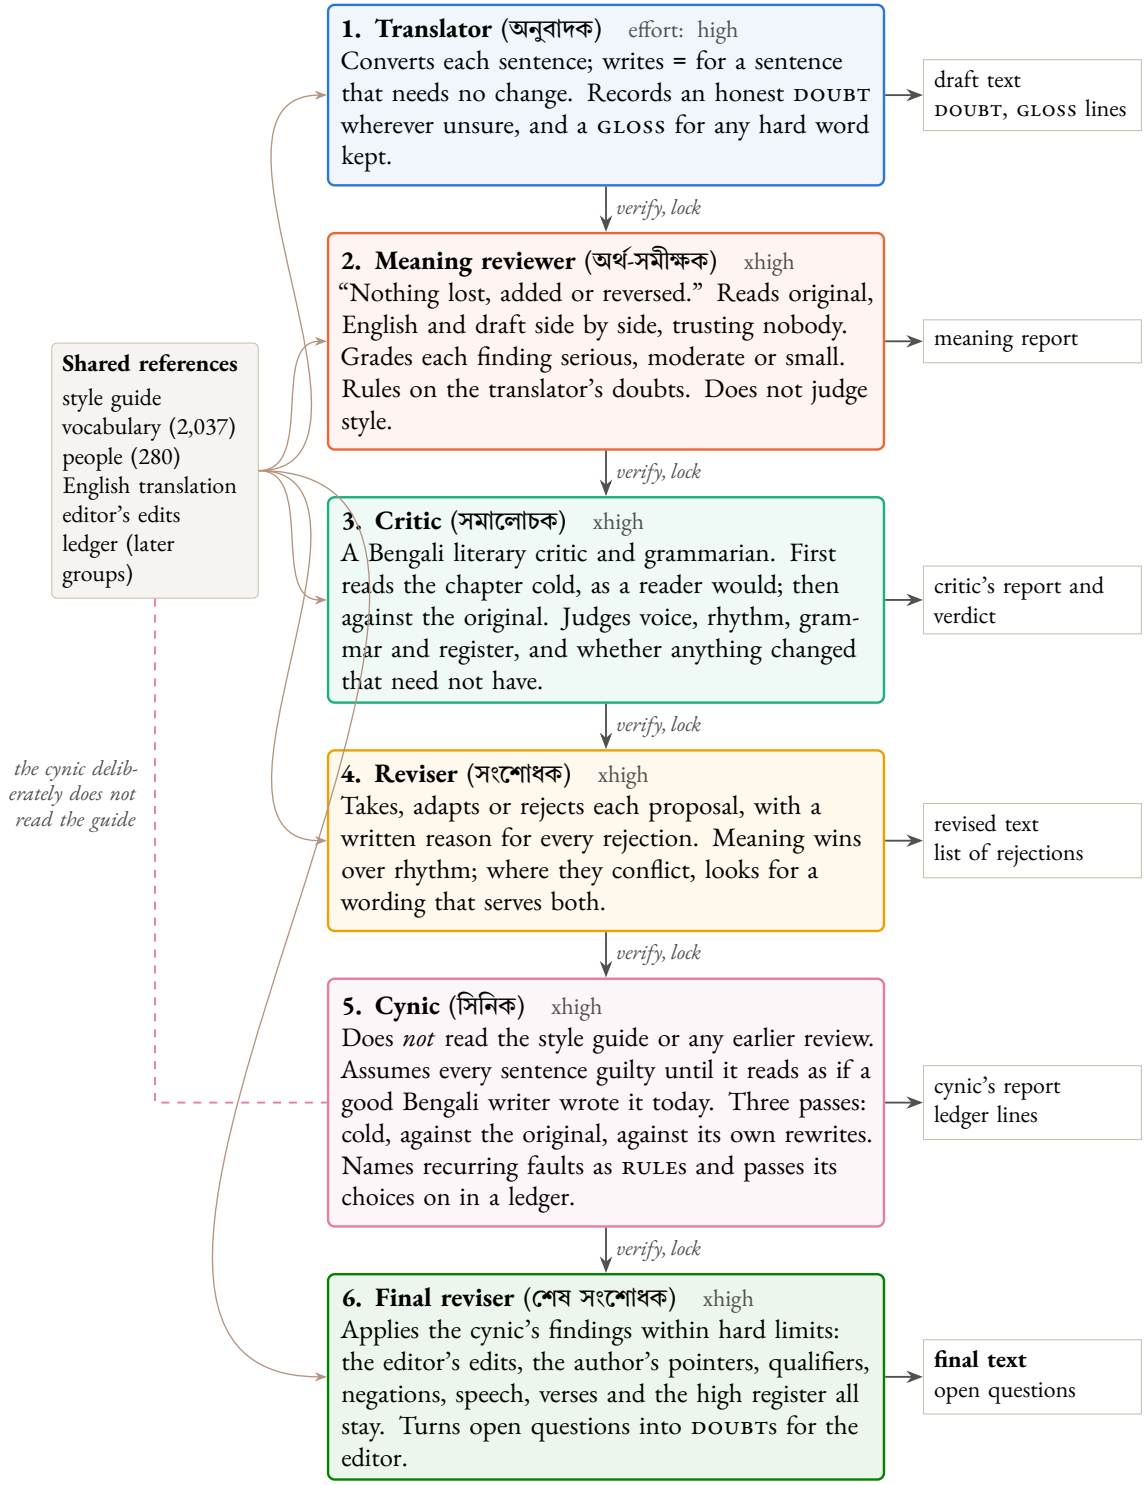


Figure 2: The six roles. Each is a separate agent working from its own brief and communicating only through files. After each stage a check program verifies the files and locks them. The cynic is kept away from the style guide on purpose, because sometimes the guide itself caused a fault and everyone who trusted it passed it.

Separate jobs. Meaning and style are judged by different reviewers. A meaning reviewer who also judges style tends to trade one for the other. A style critic who also checks facts reads too slowly to hear the prose. The meaning reviewer is told not to judge style unless it changes the meaning. The critic is told not to check omissions line by line, though it reports any meaning problem it happens to see.

Independent readings. The critic and the cynic each read the chapter *cold* before seeing the original, as a reader would. A sentence that sounds translated is noticed then, before knowledge of the original excuses it. The cynic was added after the pilot. The editor was still finding faults that every reviewer had passed, because the reviewers had judged the text against the style guide and some faults came from the guide itself. The cynic's brief says so plainly. Its only fixed points are the editor's limits (Section 2) and the editor's own corrections, given as calibration.

Objections must be sentences. Every finding must come with the full corrected sentence. "This is weak" without a better sentence is not a finding. The cynic's brief puts it this way: "Thirty real findings beat a hundred fussy ones."

Rejections must be reasoned. The revisers may reject any proposal, but each rejection is written down with its reason. Over the book the reviser rejected 144 proposals and the final reviser 349. These lists are part of the record the editor reads.

Honest doubt. The translator's brief says: "An honest doubt is a success, and a confident guess that turns out wrong is the worst outcome here." Doubts have types: meaning, who is meant, negation, word, the English differs, possible misprint, other.

Scripts never change the text. The editor ruled at the outset that the text is never changed by formula. Automatic *calit* converters over-correct. So every word of the result is an agent's judgement or the editor's own. Programs only *flag* possible problems for an agent to judge. They flag, for example: a sentence marked unchanged although it contains *sādhu* forms; a changed count of numerals, ৺ or negatives; a missing qualifier; a name form that vanished; মন্দির where the source has ঠাকুরবাটী. Programs also *verify* that a stage's files are complete. Even the one book-wide mechanical edit, modern কী for the interrogative কি, was decided occurrence by occurrence: four agents read all 1,201 occurrences in context and chose 273 to change. The program then applied exactly those 273, and asserted that nothing but the one vowel sign had changed.

6 Following one sentence

Table 3 follows sentence `p-705.2` through every stage. Mathur has asked the young Ramakrishna to take up service in the temple, and Ramakrishna states the condition he had told Hriday earlier.
Reviewers also check one another. In the heading `p-2841.0`, ঠাকুরের সাধুদের সহিত মিলন কিরূপে হয় ("How the Master met holy men"), both the meaning reviewer and the critic replaced মিলন with মেলামেশা, since মিলন today suggests union. The cynic objected that this was *over-correction*: the heading is about how the meetings came about, not about mixing with the monks, and the author's word works in *calit*. The final reviser restored মিলন. The design does not trust any single reviewer, including the last one.

7 The editor in the loop

The agents do the reading. The editor decides what the book should be (Figure 3).

The pilot. Before the full run, a pilot of 1,602 sentences from four parts of the book went through the first four stages. The editor read it. Their verdicts and examples became Section 10 of the style guide, "What the pilot taught". In the pilot the meaning reviewer found no serious faults, 3 moderate and 7 small. The critics accepted all four chapters with corrections and made 230 findings. The commonest faults were not meaning errors. They were reflex words (তেমন, সে বিষয়, তাতে dropped in without reading the sentence) and unnecessary unpacking of

Original	ঠাকুর তখন মথুরবাবুর নিকট উপস্থিত হইলেন এবং তাঁহার দ্বারা দেবালয়ে কর্ম স্বীকার করিতে অনুরুদ্ধ হইয়া পূর্বোক্ত অভিপ্রায় প্রকাশ করিলেন। English: The Master then went to Mathur and, requested by him to take service in the temple, expressed his conditional assent.
Translator	ঠাকুর তখন মথুরবাবুর কাছে গেলেন এবং **তিনি** দেবালয়ে কাজ নিতে অনুরোধ করলে ওই অভিপ্রায় প্রকাশ করলেন।
Meaning	*Who, moderate.* The subject of the previous verb is the Master, so তিনি reads as the Master, but it was Mathur who asked. The original's passive avoided this. Name the subject.
Critic	*Unclear*, found independently: as written, the Master seems to be asking himself. Proposes **মথুর তাঁকে** …
Reviser	ঠাকুর তখন মথুরবাবুর কাছে গেলেন এবং **মথুরবাবু তাঁকে** দেবালয়ে কাজ নিতে অনুরোধ করলে **সেই** অভিপ্রায় প্রকাশ করলেন।
Cynic	*Stumble.* In this chapter "the intention" has so far been Mathur's (**মথুরের এই অভিপ্রায়**), and Mathur is now the nearest subject, so সেই অভিপ্রায় reads as Mathur's. The original's পূর্বোক্ত অভিপ্রায় is the Master's condition, told to Hriday.
Final	ঠাকুর তখন মথুরবাবুর কাছে গেলেন এবং মথুরবাবু তাঁকে দেবালয়ে কাজ নিতে অনুরোধ করলে **নিজের** সেই অভিপ্রায় প্রকাশ করলেন।

Table 3: One sentence through the six stages. Two reviewers found the first fault independently; the reviser's fix opened a second ambiguity, which only a reader holding the whole chapter in mind would notice, and the cynic did.

the author's compact forms: as -আদি was opened out, প্রভৃতি rose from 8 occurrences in the original to 52 in the draft, and রকম from 24 to 101.

A blind test of the cynic. Before adding the cynic to every group, we tested it blind on one chapter of 177 sentences that had already been through the critic. Without seeing the editor's corrections, it found all six faults the editor had found in that chapter, and two of its rewrites matched the editor's word for word. It made 87 further findings. A second run at a different reasoning effort also found all six, but of the 80 sentences both runs flagged, only 23 received the same rewrite. Judgement of style varies from run to run, which is one more reason the final reviser checks the cynic.

Rulings during and after the run. Each ruling was written into the style guide and applied to finished text by hand or by a reading agent, never by search-and-replace:

- household scenes go plain;
- ইতরসাধারণ is সাধারণ মানুষ (17 finished sentences corrected by hand);
- *sādhu* grammatical forms inside mixed colloquial speech are converted (116 sentences);
- modern কী for the interrogative (273 changes);
- converted *sādhu* habitual verbs take the standard *-tām* form;
- দান is kept where the book speaks of giving instruction, initiation or darśan.

Where several readings were possible, the editor decided on a review page that showed each full sentence with the alternatives.

The editor's own sentences. 402 sentences now carry wording fixed by the editor: about 150 written by the editor or corrected by hand on the editor's ruling, and 261 changed under the কী ruling. Each edit is logged with the sentence before and after. Every later stage is told that the editor's wording stands, and an audit confirms after every change that all 402 still read exactly as the editor wrote them.

The ledger. Each cynic appended up to ten lines to a ledger for its part of the book. These record choices that later chapters must keep the same, for example "পূর্বোক্ত X → যে X-এর কথা আগে বলা হয়েছে" or "মাতা, ভ্রাতা in

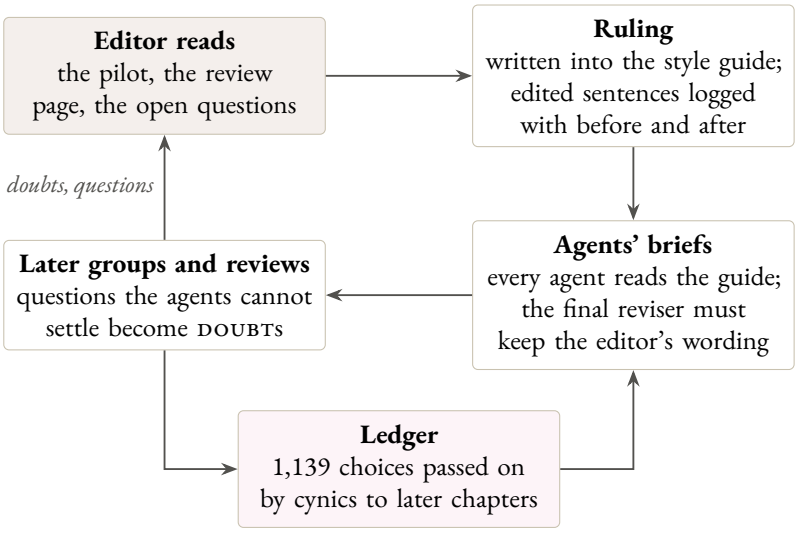


Figure 3: How rulings propagate. A ruling made once reaches every later group through the style guide. A sentence the editor has written is protected from every later stage and checked by an audit. The cynics' ledger keeps later chapters consistent with earlier choices.

family scenes → মা, ভাই; মাতৃভক্ত, জননী stay". Translators and revisers of later groups read the ledger, so the same fault is not made again. By the end it held 1,139 lines.

8 Guarding what must not change

The project kept all its state on disk. Nothing depended on a conversation staying alive, which mattered: the run was interrupted by usage limits and resumed by different sessions. Each stage wrote only to its own folder. A stage was finished only when a check program had confirmed that:

- every sentence id is present, in order, exactly once;
- there are no empty or stray lines;
- the review files exist and have the required form;
- at the final stage, footnote marks and ৺ match the original sentence by sentence.

The program then marked the stage done and locked its files read-only. A stage interrupted by a usage limit resumed from its last saved batch; no finished work was lost or redone.

After the run, a final audit compared the whole book with the original (Table 4).

	Original	Final	
Footnote marks	510	510	equal
The sign ৺	305	305	equal
Numerals	719	719	equal
যেন ("as if")	364	375	never fewer; +11 where যাহাতে ... না ("lest") became যেন ... না
প্রায় ("almost")	265	266	never fewer
বোধ হয় ("perhaps")	163	170	never fewer; বোধ হইয়া থাকে became বোধ হয়
হয়তো ("maybe")	60	60	equal
Editor's edited sentences	—	402	all as the editor wrote them

Table 4: The final audit. Every place where a qualifier was added was read by hand.

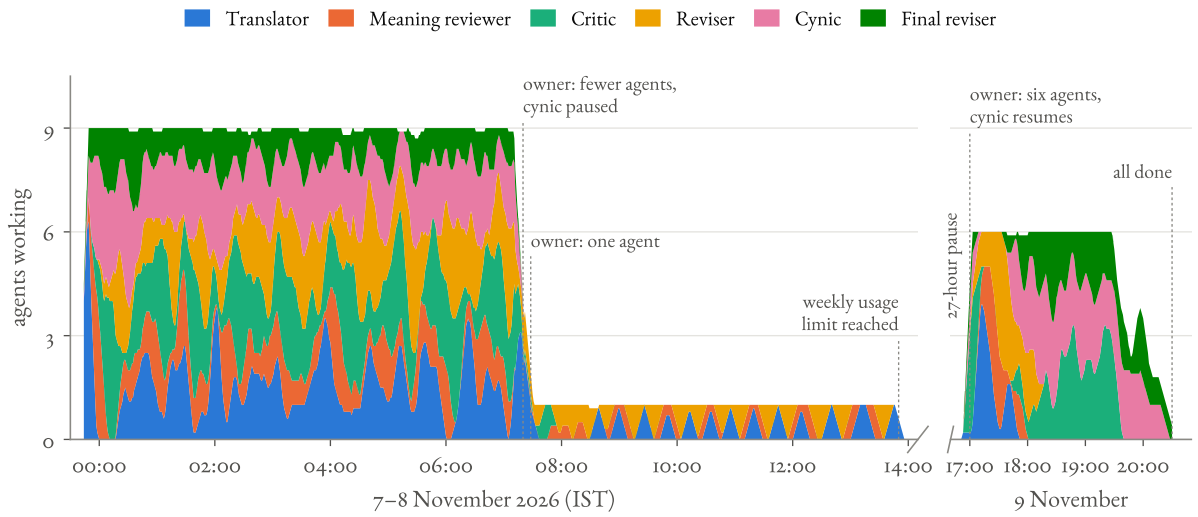


Figure 4: Agents at work over the full run, by role (ten-minute averages). The gap is the pause at the weekly usage limit. The pilot, earlier on 7 November, is not shown.

9 Running it

The full run started at 23:44 on 7 November 2026 and ended at 20:30 on 9 November (India time). Agents were launched by workflow scripts in Claude Code. Each agent was given its brief and group, and the next stage of a group started as soon as the previous one was verified. Overnight nine agents worked at once. In the morning the editor first reduced them and paused the cynic, then cut them to one agent as the weekly usage allowance ran low. The run stopped for 27 hours at the weekly limit and resumed with six agents (Figure 4). Actual running time was about 17.7 hours.

Role	Agent runs	Agent-hours	Output tokens (millions)
Translator	118	16.2	7.20
Meaning reviewer	110	10.5	4.47
Critic	126	18.8	7.59
Reviser	109	16.0	6.91
Cynic	119	19.7	8.32
Final reviser	118	11.1	4.56
Total	700	92.3	39.04

Table 5: Measured cost of the full run, summed from the agents' own logs. Runs include restarts after interruptions and the critic's quick exits for the groups it reviewed later. The first four stages of the ten pilot groups ran earlier and took about 3.7 agent-hours more. Besides their output, the agents read about 1.43 billion tokens of context, most of it from cache: the style guide, briefs and texts read again at every step. All agents were Claude Opus 5.5.

Two features of the cost stand out. Reviewing cost more than translating: the five reviewing and revising roles used 76 of the 92 agent-hours. And most of what the agents read was not the book. It was the style guide, the briefs, the people and the vocabulary, read again by every agent. That is the price of consistency, and it is exactly what a one-line prompt leaves out.

10 Results

Step	Sentences changed	Words changed
First draft, against the original	13,828 (82.1%)	36.4%
Reviser, against the draft	2,446 (14.5%)	1.5%
Final reviser, against the revised text	2,599 (15.4%)	2.4%
Editor, after the run	393 (2.3%)	0.2%
Final text, against the first draft	4,642 (27.6%)	
Final text, against the original	14,080 (83.6%)	38.1%

Table 6: What each step changed. "Words changed" is the share of the earlier version's words not kept, in order, in the later version. The 2,758 sentences that are identical to the original are mostly verses, speech, headings and sentences already in a form modern readers use.

Review changed more than one sentence in four of the first draft, and the draft was already the work of a well-prepared translator (Table 6). In words the changes are small: a few per cent. But small changes are where meaning lives. Every fault in Table 2 was fixed by a change of one or two words.

Reviewer	Findings	By kind
Meaning reviewer	153	2 serious, 24 moderate, 127 small. Wrong sense of a word 75; wrong person 29; something omitted 14; level of respect 8; fixed text not kept 7; something added 5; tense 3; negation 3; other 9. Also 1,315 rulings on the translator's doubts.
Critic	1,263	unclear 290; too literal 275; grammar 169; rhythm 127; voice 118; monotonous 116; changed too much 87; spelling 45; *sādhu* and *calit* mixed 35; sentence division 1. Every group accepted with corrections; none sent back.
Cynic	2,611	stiff (still *sādhu* in shape) 1,239; a reader stumbles 912; reflex word 312; over-correction 80; meaning moved 68. Also 125 questions for the editor and 633 named patterns (RULEs).
Rejected	493	by the reviser 144 (101 from the critic, 43 from the meaning reviewer); by the final reviser 349. Each with a written reason.

Table 7: What the reviewers found. 3,460 sentences (20.5%) drew at least one finding.

The pattern in Table 7 is worth stating plainly. Errors of meaning were rare: the meaning reviewer found one in about every 110 sentences. But they were real, and two would have told readers something false. Faults of style were common, with nearly 3,900 findings. Even after the critic had passed a chapter and the reviser had rewritten it, the cynic still flagged 2,531 sentences, roughly one in seven. Without the cynic, the book would have been accurate and stiff.

Open questions. The translators raised 1,337 doubts. The reviewers settled most of them, and 518 remain open for the editor, on 503 sentences. Of these, 215 are possible misprints in the source. An example is `p-17.0`, where যাহা— looks like a misprint for যথা— ("namely"), which the author writes elsewhere. 111 note places where the English differs from the Bengali in a way that might matter; 27 concern a word, 16 who is meant, 15 a meaning, and 134 other matters. Some need the printed book: one sentence promises a family tree that the digital text does not contain. Others need a decision only the editor can make.

Notes for the reader. The final text carries 2,678 sentence-specific glosses. In addition, single-sense KEEP words from the vocabulary are underlined automatically. A glossary at the end of the book collects them.

11 What the reader gets

The Saral reader (Figure 5, top) shows the modern text. A click on any sentence opens the original sentence beneath it, and a switch reverses the arrangement: the book in the original, with the modern sentence a click away. Words with a dotted underline show their meaning when clicked. The reader can also change the text size and use a full table of contents. On a first visit a short note, in Bengali and in English, explains why this version exists, says that it is a bridge to the original and not a substitute, and states that AI was used and how the text was checked.

For reviewers there is a second page (Figure 5, bottom). It shows original, modern text and English side by side, sentence by sentence, with every open question marked. A filter shows only those sentences. Beside each question are buttons for "this is right" and "this must change" and a box for comments, and the reviewer's decisions can be exported as one file. Both pages pass 45 automated checks in a headless browser.

12 Limitations

- **The text was produced by AI.** Every role was played by the same model family. Different briefs, cold readings, the cynic's ignorance of the guide and the editor's rulings reduce the risk of shared blind spots, but they cannot remove it.
- **518 questions are open.** Of these, 215 are suspected misprints in the source. The digital text has not yet been collated against a printed edition.
- **No independent evaluation yet.** No outside Bengali reader has yet rated a sample of the text blind, against the original or against a human rendering. Such an evaluation is the obvious next step, and the side-by-side reader was built partly to make it easy.
- **Single runs.** Each group was passed once by each role. The pilot comparison of two reasoning efforts showed that style judgements vary between runs, and we have not measured that variation at scale.
- **The register is a choice.** Level A is one reasonable target. Readers who want plainer Bengali, or who prefer the original's cadence throughout, will each find places to disagree.
- **Cost.** Ninety-two agent-hours and over a billion tokens of context are not trivial. We report them so that others can judge whether the method fits their project.

13 Conclusion

Asking a model to "convert this book into modern Bengali" would have produced fluent text quickly, and that text would have contained, unannounced, every kind of fault in Table 2. What made this edition fit to publish was old editorial practice carried out at the scale of a whole book. A clear target was set by the editor's verdicts. Reference material was ready before the first sentence. Meaning and style had separate readers, and one reader ignored the rules on purpose. Every rejection had a written reason. The editor's rulings reached every page, and machine checks guarded what must not change. AI agents made this affordable. They did not make it unnecessary.

The Saral Līlāprasaṅga does not replace the original. It is a bridge to it, and every sentence keeps the original one click away. Corrections from readers are welcome.

Statement on the use of AI

The modern-Bengali text was produced by AI agents (Claude Opus 5.5, Anthropic, run through Claude Code) under the direction of the editor, as described in this report. This report was drafted with the help of the same AI system from the project's files and logs. Every figure in it was computed from those files and logs. The author is responsible for its content.



Figure 5: Top: the Saral reader, with sentence `p-705.2` clicked to show the original beneath it. Bottom: the side-by-side reader for reviewers, with original, modern text and English, and an open question awaiting a decision: the English mentions Hriday's appointment, which this Bengali sentence does not, and the Bengali was followed.

Data availability

The published edition is at `https://cs.rkmvu.ac.in/~tamal/saral-lilaprasanga.html`. The agents' briefs, the style guide, the vocabulary, the list of people, the check programs and the side-by-side reader are available from the author on request.

References

[1] Swami Saradananda. *Śrīśrīrāmakṛṣṇa-Līlāprasaṅga*. Five parts, written 1909–1919. Kolkata: Udbodhan.

[2] Swami Saradananda. *Sri Ramakrishna the Great Master*. Translated by Swami Jagadananda. Madras: Sri Ramakrishna Math, 1952.

[3] Charles A. Ferguson. "Diglossia." *Word* 15, no. 2 (1959): 325–340.

[4] Suniti Kumar Chatterji. *The Origin and Development of the Bengali Language*. Calcutta: Calcutta University Press, 1926.