

अलि अलि पिरौ
अलि अलि नुनिलो
अलि अलि गलिया

AI & TECHNOLOGY / AI & TECHNOLOGY

Can AI Understand Nepal's 124 Mother Tongues?

Nepal's language diversity presents a test for artificial intelligence: whose words are recorded, whose speech is understood and who gets to decide?

OCTOBER 8, 2026



Can AI Understand Nepal's 124 Mother Tongues?

Nepal's language diversity presents a test for artificial intelligence: whose words are recorded, whose speech is understood and who gets to decide?

By Nepal Digest - October 8, 2026 - Nepal

Ask an artificial intelligence assistant a question in English and it may respond with a polished paragraph. Ask it in Nepali and the answer may still sound confident. Switch to a smaller language spoken in a Himalayan valley or a community of the Terai, however, and the experience may change: the model may misread a word, choose a neighbouring language or produce a fluent-looking response that means little to the person who asked.

This is not merely a technical inconvenience. Language determines who can ask for information comfortably, whose knowledge appears in digital systems and which forms of culture are treated as understandable.

For Nepal, the question is much bigger than whether a chatbot can write grammatical Nepali. It is whether a digital future built around machine language can serve a country where hundreds of communities speak in different ways.

The number that changes the question

Nepal's National Population and Housing Census 2021 recorded 124 reported mother tongues. The National Statistics Office's language thematic report gives Nepali the largest share, approximately 44.9 percent of the population identifying it as their mother tongue. It also documents the enormous variety beyond the most widely spoken languages.

These figures come from self-reported census data and administrative linguistic classification. They are not an inventory of the languages an AI model understands, nor a measurement of the number of fully separate dialects or writing systems. A language may appear in census results yet have little readily available digital text. Others are spoken extensively in homes and communities while appearing rarely in online documents.

That mismatch is central to the AI problem. Machines do not acquire language ability simply because a government recognises a language. They need suitable data, evaluation methods, technical tools and sustained collaboration.

What does it mean for a machine to ‘understand’?

The word ‘understand’ carries more promise than most AI testing can support. A model might recognise a script but struggle to answer a detailed question. A speech system might transcribe a standard accent yet perform poorly with another variety. A translation program could preserve literal vocabulary while destroying irony, respect or context.

It helps to separate different tasks: recognising speech, converting speech to text, reading characters, translating between languages, searching for information and generating an appropriate response. Proficiency on one task does not guarantee competence on another.

A system that occasionally produces an impressive paragraph is not automatically reliable enough for health advice, legal interpretation or classroom teaching. Reliable performance needs tests with representative speakers, contexts and dialects; it also needs an honest account of errors.

No broad assertion that one named commercial model is ‘best for Nepal’s languages’ should be published without current, reproducible benchmark evidence. Technologies change rapidly, and vendor demonstrations are often selective.

The internet does not contain every language equally

Much contemporary language AI is developed using large collections of digital text and other data. Languages with abundant books, web pages, subtitles and labelled recordings have a different starting position from those with limited digital resources.

For many smaller languages, useful speech or text exists but may not have been transcribed, digitised, standardised or made available under appropriate conditions. The absence of a large dataset does not indicate a language is simple, undeveloped or lacking expressive power. It reflects what has been documented, funded and shared.

A community can have an extensive oral tradition and only a small searchable corpus. A language may have more than one accepted spelling convention or use scripts not well supported by common keyboards and fonts. One speaker's vocabulary may be unfamiliar in a neighbouring district.

AI systems trained without careful local evaluation can flatten these differences, returning the dominant form of a language while treating others as mistakes.

The problem is partly engineering and partly a decision about whose speech counts.

A child in a multilingual classroom

The consequences can begin long before an adult uses a chatbot at work. Students who speak one language at home may meet different classroom languages at school. A digital tutor trained only on major languages may be accessible to some pupils and frustrating to others.

A 2024 study commissioned by UNICEF Nepal with CEHRD, the Language Commission, UNESCO and other partners examined how language policies and mediums of instruction influence students' experiences and outcomes.

That report did not establish that generative AI can solve multilingual schooling. It demonstrates why language is a material educational issue rather than a decorative part of cultural identity.

If AI tools are introduced into classrooms, teachers need to know what they can and cannot do. A language model might help prepare vocabulary practice or draft a reading passage for review. It can also invent a definition, mishandle a culturally specific expression or teach a form that speakers consider wrong.

A fluent interface should not be mistaken for a qualified local teacher.

Whose stories become training data?

Suppose researchers want to create a speech-recognition system for a language with few public recordings. Gathering community interviews, songs, oral histories and conversations might improve the technology. But those materials belong to people and traditions, not merely to the team building the dataset.

A recording provided for linguistic preservation has not necessarily been authorised for unlimited commercial machine-learning use. Speakers may discuss sensitive family, sacred or community knowledge. Some may not wish particular materials to be made public or repurposed.

Informed consent therefore has to be more than a generic checkbox. Contributors need to understand who will use recordings, whether the model or data may be shared, how correction and withdrawal requests will be handled, and what benefit the community might receive.

Governance matters especially when community members cannot independently audit the systems built from their words. Data collection can otherwise reproduce an old pattern: knowledge leaves a community and its economic or scholarly value accumulates elsewhere.

A language technology project should begin by asking whom it serves and who controls it.

Why ‘supporting a language’ is not a single feature

A useful tool for one community may be an audio dictionary that helps young people hear elders pronounce words. Another may need accessible keyboard input, optical recognition for printed texts, search across local publications or a tool for transcribing oral histories.

A universal chatbot may be the least urgent need.

Different use cases demand different error tolerances. A creative writing aid can withstand uncertainty if the user knows the limitations. A public-health translation or emergency warning cannot. Building the latter requires thorough testing, human review and institutional accountability.

Small, carefully evaluated applications may therefore be more valuable than a grand promise to make a foundation model ‘fluent in every Nepali language’.

A sensible development sequence begins with the community’s needs, builds usable resources, establishes baselines and publishes evidence of accuracy. The model comes after the problem, not before it.

The risk of linguistic flattening

In a 2025 UNESCO discussion of AI and linguistic flattening, the author examines how generated language can develop repetitive patterns and reduce diversity of expression. The observation is broad, not a field study of Nepal's smaller languages, but it raises an important question.

When a system makes one way of speaking feel normal and many others unusual, it can influence choices. Users may adjust their questions to fit the machine, change spellings that the tool fails to recognise or switch to English to avoid poor results.

Those adaptations may make the tool easier to use. Over time they can also push local expression out of digital spaces where it deserves to be present.

AI cannot single-handedly preserve a language. Family use, schooling, literature, community institutions and intergenerational transmission remain indispensable. But poorly designed technology can make a language less visible precisely when digital communication occupies more of daily life.

Who should build the tools?

Universities, community organisations, local developers, public institutions and international researchers can contribute different capabilities. Academic institutions can develop evaluation

methods. Governments can support public-interest resources. Local speakers can establish what counts as an appropriate translation or socially acceptable utterance.

No one group should claim complete authority over a living language. Diversity exists within communities as well as between them; an evaluation programme that tests only one elite register may exclude people who speak differently.

A stronger system would be transparent about supported dialects, data sources and limitations. It would compensate contributors fairly, offer tools usable on low-cost devices and make corrections possible after deployment.

The benefit might extend well beyond culture: clearer access to public information, assistive technology, learning resources and services for people who do not find English or standard Nepali comfortable.

But every such claim needs testing with actual users.

The question behind the next prompt

A person using a digital assistant should not have to abandon a mother tongue to receive a competent answer. That is a worthwhile ambition. It is also one for which technological optimism is not evidence of achievement.

Nepal's 124 reported mother tongues represent people, not 124 items on a product roadmap. Some need documentation, others demand better tools, and every one exists in a living relationship with its speakers.

The measure of progress is not whether a model can impress a researcher in a demonstration.

It is whether people can speak as themselves—and still be understood.

Reporting note: This feature discusses technical limitations and governance risks; it does not rank current AI products or claim original performance testing of Nepal's languages.

Originally published by Nepal Digest - <https://nepaldigest.com/can-ai-understand-nepals-124-mother-tongues/>