

IIT ROORKEE · DEPARTMENT OF ELECTRICAL ENGINEERING · PSQUARE LAB

AI, and *AI for Bharat*. What it means for your career.

A few rules that I believe will make you thank yourself in 2036.

Parikshit Pareek · Assistant Professor, EE, IIT Roorkee

Parts of this slide deck were prepared using Fable 5

WHO IS TALKING TO YOU

Parikshit Pareek — the *journey* so far.



Point of this slide: none of these doors needed a famous school to open the first one. They needed maths, stubbornness, and one good problem at a time.

START HERE

When someone says "AI", they mean *this*.

ChatGPT. Gemini. Claude. Meta AI in WhatsApp. You type, it talks back, and it sounds like it knows everything.

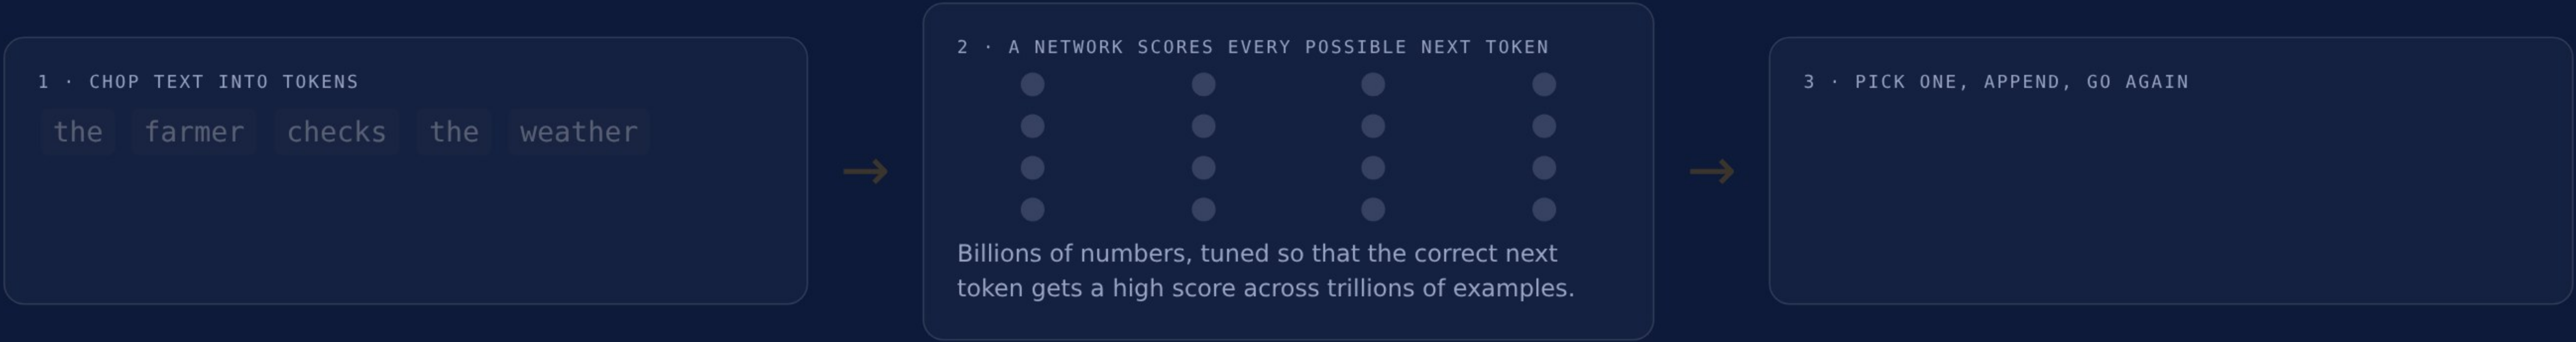
These are *large language models*. They are one kind of AI, and the kind that will touch your career most.

So let's open them up!

A language model is a machine that has read a very large amount of text and learnt to guess **the next word**.

HOW IT WORKS, IN ONE LOOP

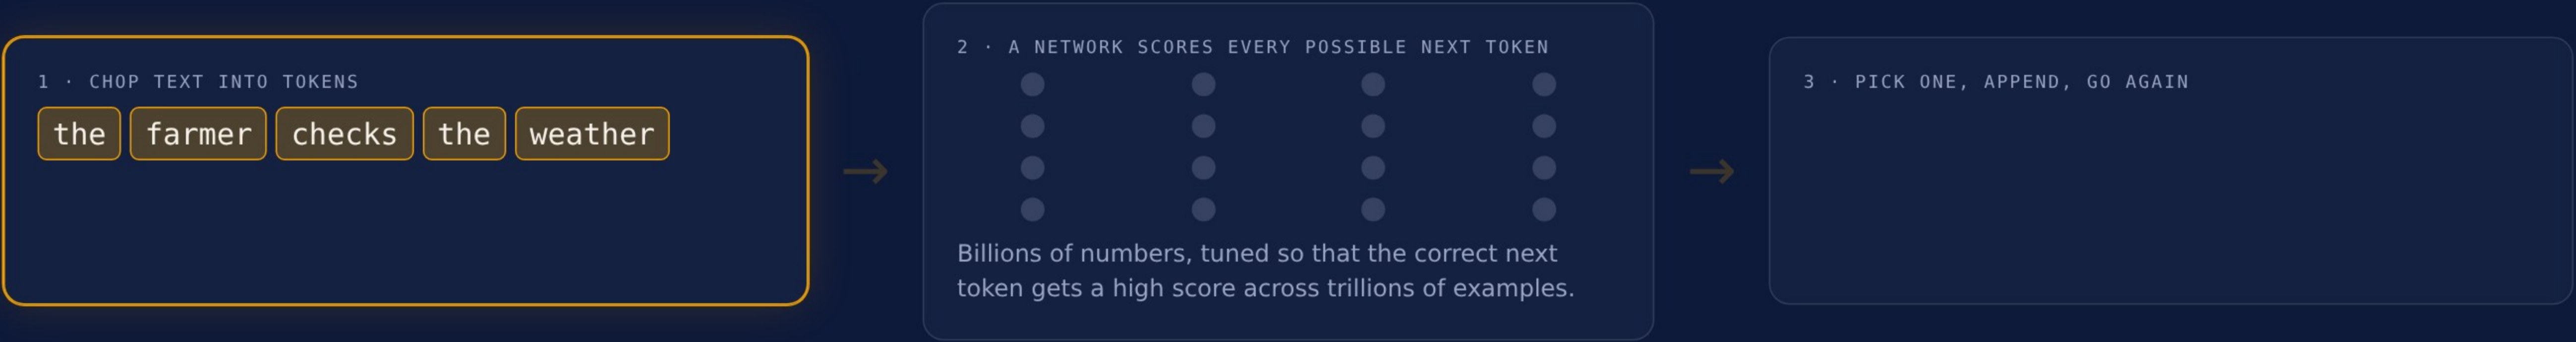
Chop. Score. Pick. *Repeat.*



Press → to run one step of the loop.

HOW IT WORKS, IN ONE LOOP

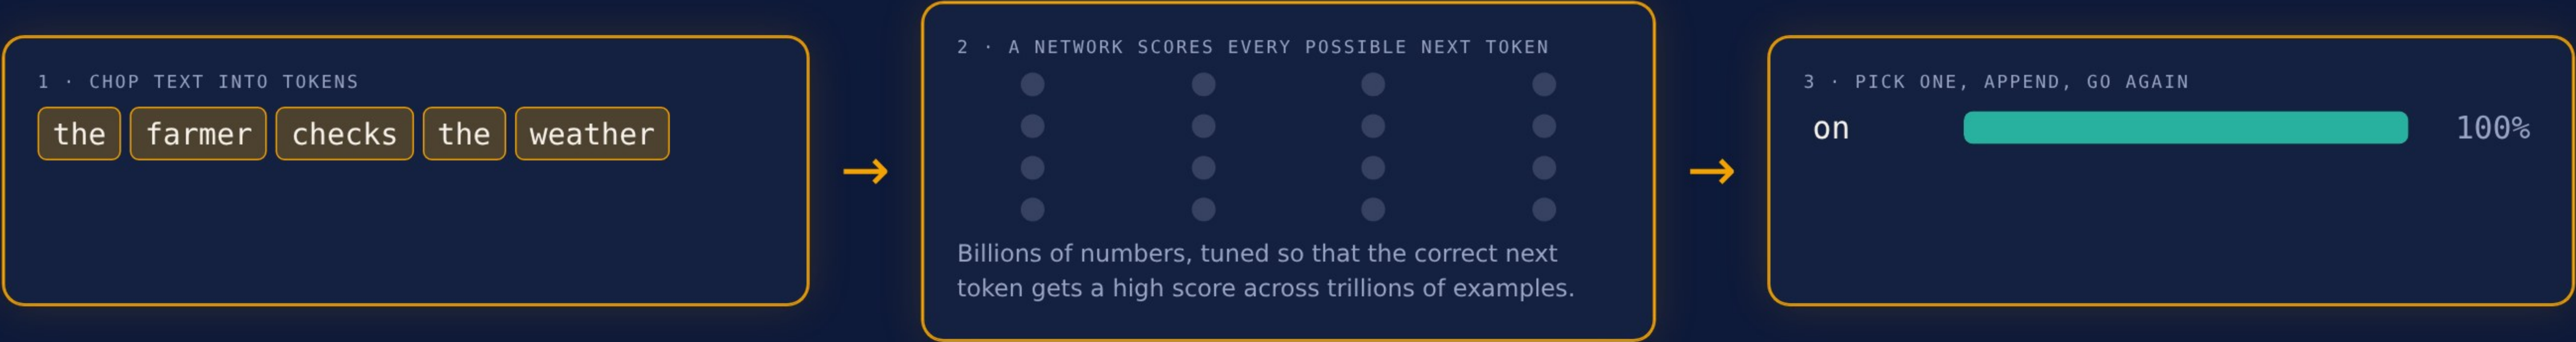
Chop. Score. Pick. *Repeat.*



Step 1: the sentence is chopped into tokens.

HOW IT WORKS, IN ONE LOOP

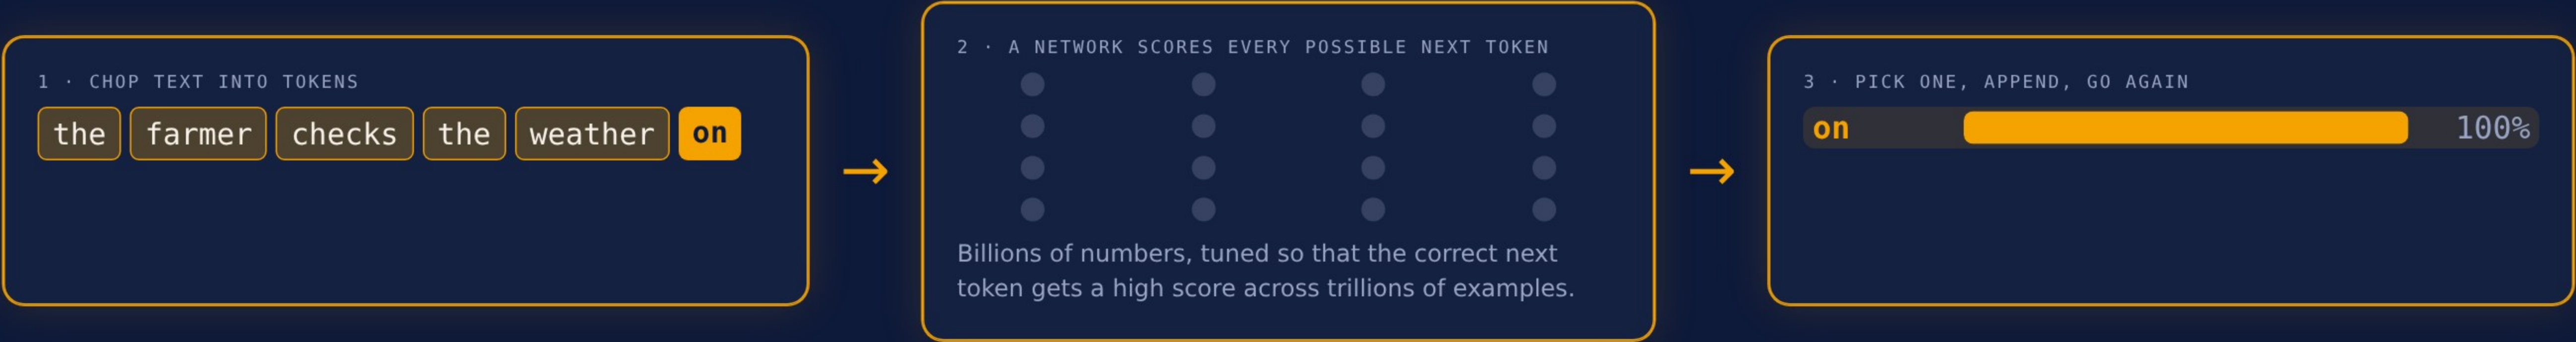
Chop. Score. Pick. *Repeat.*



Step 2: every candidate next token gets a score. Scores become probabilities.

HOW IT WORKS, IN ONE LOOP

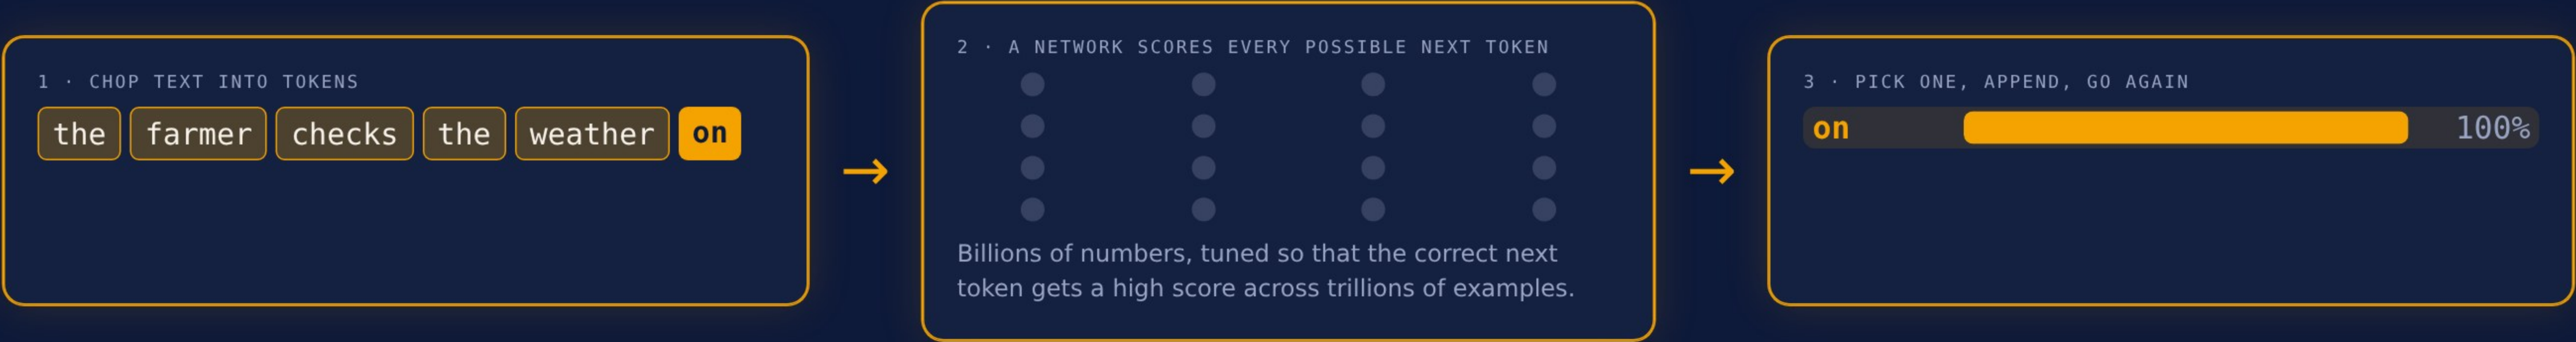
Chop. Score. Pick. *Repeat.*



Step 3: pick one, usually a likely one, and stick it on the end.

HOW IT WORKS, IN ONE LOOP

Chop. Score. Pick. *Repeat.*



That is the whole machine. Everything else is scale and engineering. Press → again to try it yourself.

LIVE EXPERIMENT · A REAL (TINY) LANGUAGE MODEL RUNNING IN THIS PAGE

Watch it guess the next word.

the farmer in roorkee explains the **problem**

TOP CANDIDATES AND THEIR PROBABILITIES

again 100.0%

Prediction uses: last two words. 1 possible next words seen in training text.

"the farmer in ..."



Predict next word

Auto-run

Start over

Randomness (temperature) 1.0

$P(\text{next word} \mid \text{last two words}) = \text{count}(\text{last two words} + \text{next word}) \div \text{count}(\text{last two words})$

This model has read only 560 words. ChatGPT-class models have read trillions, and use a neural network instead of a counting table. The idea is the same: predict, pick, repeat.

THREE THINGS YOU JUST SAW

It predicts. It does not know.

The model never checked whether a sentence is true. It only checked whether the words are *likely*. Truth is a side effect of reading a lot of mostly-true text.

It is fluent even when wrong.

Turn up randomness and the sentences stay grammatical but drift into nonsense. Big models do this too, just more rarely. This is called a **hallucination**: a confident, fluent, false answer.

It learns what it reads.

Our model knew "Roorkee" and "monsoon" because they were in its text. It knew nothing else. A model that has read little about Bharat will be weak on Bharat. Hold that thought.

WHY DOES IT FEEL SO SMART, THEN?

10^{12}

words of text, roughly, read during training by a frontier model. A human reading a book a week for 80 years reads about 4×10^8 words.

Scale makes the guesses astonishingly good. It does not change what the machine is doing.

10^{11}

adjustable numbers ("parameters") inside the network that store the patterns. Our demo had a few thousand counts.

Orders of magnitude only. Exact figures for the newest models are not public.

PART TWO

Your brain is the *asset*. Do not outsource its workout.


The single biggest career risk for your generation is not "AI will take my job". It is "I never built the skills because AI did my homework".

HOW SKILL IS BUILT

Struggle is not a bug. It is the *mechanism*.

Learning science calls it **desirable difficulty**: you remember and understand what you had to work to produce. Retrieving an answer from your own head strengthens the memory. Reading the answer does not.

You cannot pay someone to lift weights for you and expect your arms to grow. Cognition works the same way.



Illustrative shape, not measured data. The early-years overlap is the trap: both paths feel fine until one of them stops.

THIS IS NOT A TEACHER'S OPINION. IT IS BEING MEASURED.

Four studies, 2025.

PNAS · 2025 · RANDOMISED TRIAL, ~1,000 STUDENTS

Generative AI without guardrails can harm learning: evidence from high school mathematics

Bastani, Bastani, Sungu, Ge, Kabakcı, Mariman · Wharton / Penn

High-school students in Classes 9–11 a guarded tutor, or nothing. Then everyone practised maths with a ChatGPT-like tool, sat a closed-book exam.

With plain ChatGPT: better in practice, worse on the exam than students with no AI. Guarded tutor: no harm, no gain.

MIT MEDIA LAB · ARXIV 2506.08872 · EEG STUDY

Your Brain on ChatGPT: accumulation of cognitive debt when using an AI assistant for essay writing

Kosmyna, Hauptmann, Yuan, Situ, Liao, Beresnitzky, Braunstein, Maes

54 participants wrote essays over four months with ChatGPT, with search, or with brain only, while brain activity was recorded.

Weakest brain connectivity in the ChatGPT group; poorer recall of their own essays. Preprint; sample is small and critiques exist.

SOCIETIES · 2025 · SURVEY + INTERVIEWS, 666 PEOPLE

AI tools in society: impacts on cognitive offloading and the future of critical thinking

Gerlich · SBS Swiss Business School

Measured AI use, "cognitive offloading" thinking scores across age groups. (letting the tool think for you), and critical-

Heavier AI use went with lower critical-thinking scores, strongest in the youngest group. Correlation, not proof of cause.

CHI · 2025 · SURVEY OF 319 KNOWLEDGE WORKERS

The impact of generative AI on critical thinking: self-reported reductions in cognitive effort and confidence effects

Lee, Sarkar, Tankelevitch, Drosos, Rintel, Banks, Wilson · Microsoft Research / CMU

Professionals described 900+ real tasks for themselves. done with AI and how much they thought

The more they trusted the AI, the less they checked it. Critical thinking shifted from doing the work to verifying it, when it happened at all.

Read them yourself: pnas.org/doi/10.1073/pnas.2422633122 · arxiv.org/abs/2506.08872 · doi.org/10.3390/soc15010006 · Lee et al., CHI 2025. One rigorous finding beats four headlines: the PNAS trial is the one to remember.

THE DIVIDING LINE

Offload the *chores*. Keep the *thinking*.

Never outsource

- Your first attempt at any problem
- Writing to figure out what you think
- Mental maths and estimation ("is $2,300 \times 48$ about 100k?")
- Core facts and vocabulary of a subject (you cannot reason with words you do not own)
- Reading a whole chapter, not its summary
- Deciding what is true

Outsource freely

- Formatting, boilerplate, repetitive code
- Finding and translating sources
- A second opinion *after* your own attempt
- Generating practice questions for yourself
- Explaining one concept five different ways
- Tidying up grammar once the idea is yours

A RULE YOU CAN ACTUALLY FOLLOW

The *Attempt-First* protocol.

STEP 1

Attempt

Close the chatbot. Write your answer, even a bad one. Ten minutes, no help.

STEP 2

Ask

Now ask the AI. Paste *your* attempt and ask where it is wrong, not for the answer.

STEP 3

Compare

Find the exact line where you diverged. That line is the lesson. Everything else you already knew.

STEP 4

Explain back

Shut the screen and explain the solution to a friend or a wall. If you cannot, you have not learnt it yet.

The order is the whole point. Step 2 before Step 1 is how you end up with a degree and no skill.

USING IT LIKE A PROFESSIONAL

Good prompts come from *good thinking*, not the other way round.

- **Be specific.** "Explain Newton's third law using a cricket ball hitting a bat, for a Class 10 student" beats "explain physics".
- **Make it quiz you.** "Ask me 5 questions on this chapter, one at a time, and grade my answers."
- **Ask for the disagreement.** "What would a strict examiner mark wrong in this?"
- **Demand sources, then open them.** Models invent references. Click before you trust.
- **Never paste an answer you cannot defend.** In an exam, a viva, an interview: you are the one in the room.

TRY THIS TONIGHT

"Here is my solution to [problem]. Do not give me the correct answer. Tell me the first step where my reasoning breaks, and ask me a question that would help me find it myself."

PART THREE

AI takes *tasks*, not jobs.

A job is a bundle of tasks. AI removes some tasks from the bundle, makes others cheap, and makes a few far more valuable. Your career is deciding which tasks you will be known for.

EXAMPLE: A DOCTOR IN 2036



**The premium
moves to what AI
cannot be *held*
responsible for.**

- Framing the problem correctly
- Checking that an answer is actually right
- Deep domain knowledge to spot the subtle error
- Taking responsibility for the decision
- Explaining it to a human who is scared

Replace "doctor" with engineer, lawyer, civil servant, farmer, founder. The shape is the same.

WHAT CHANGES, FIELD BY FIELD

Every path you are considering *still exists*. The entry ticket changed.

Engineering

Writing code is cheap. Knowing *which* system to build, why it fails, and how to verify it is the job. Maths and physics matter more, not less.

Medicine

AI reads the scan. The doctor decides, under uncertainty, with a family watching. Biology depth plus judgement.

Law & civil services

Drafting and search get automated. Argument, ethics, and knowing the ground reality of a district do not.

Agriculture & rural enterprise

Weather, price, pest and soil advice in your own language on a ₹8,000 phone. The farmer who can use it out-earns the one who cannot.

Design, media, teaching

Generation is free; taste and trust are scarce. The teacher who uses AI to give each child a different exercise wins.

Building AI itself

Linear algebra, probability, calculus, programming, and a domain you care about. This is the path we will come back to.

PART FOUR

AI for *Bharat* will be built by whoever shows up.

Models learn what they are shown. If Bharat's languages, problems, and people are not in the data and not on the team, the model will not serve Bharat. That is not a prediction. That is arithmetic.

THE GAP, IN NUMBERS

22

scheduled languages. Over 1.4 billion people. A few hundred million speak English comfortably; most do not.

<1%

of text on the public internet, roughly, is in Hindi, and the other Indian languages are far smaller. English is around half.

Web-content surveys vary by method; the order of magnitude is what matters.

Remember the demo: a model that has barely read a language is guessing blind in that language.

PROOF IT IS NEEDED, AND THAT IT CAN BE DONE HERE: SARVAM AI, BENGALURU

Global models made Bharat pay **4× more** per word. Sarvam fixed the plumbing.

English, any tokenizer



Hindi/Tamil/Bengali, global tokenizers



Same languages, Sarvam-1 tokenizer



- **The problem.** English-first tokenizers know English words whole and shred Indian scripts, so every Indian-language user pays more and gets less.
- **What they did.** Own tokenizer for Indian scripts, ~2 trillion tokens of curated Indian-language text, a 2-billion-parameter model trained from scratch: Sarvam-1.
- **Result.** Beat much larger global models on Indian-language tests, several times faster. Small, because the plumbing was right.
- **Since then.** Sarvam-M (2025); in 2026 open 30B and 105B models with a tokenizer for all 22 scheduled languages, plus speech, translation and document models.

Founded by Indian researchers, solving a problem the giants had not bothered with. That is the template.

Tokens per word ("fertility"), as reported by Sarvam for Sarvam-1, Oct 2024. Bars drawn to the midpoint of each range. More tokens = slower, costlier, and a model that sees your language in fragments.

PROBLEMS ONLY WE WILL SOLVE, BECAUSE ONLY WE HAVE THEM

Voice, not text

Hundreds of millions of first-time internet users type poorly and speak fluently, often in a mix of languages. AI for Bharat must listen.

Public infrastructure

Aadhaar, UPI, ONDC: India builds digital public goods at population scale. Bhashini (language translation) and the IndiaAI Mission are the same idea applied to AI.

Frugal compute

Models that run on a cheap phone with patchy network, in a clinic without a data centre. Efficiency is a research field, not a compromise.

Agriculture & energy

Monsoon forecasts, crop disease from a photo, power grids with rooftop solar and batteries. Physics plus AI, both needed.

Health at the last mile

One doctor per thousand people. An assistant that helps an ASHA worker in Garhwali is worth more than one that writes English poetry.

Education in mother tongue

A patient tutor for every child, in every language, that strengthens thinking instead of replacing it. You are the first users. You could be the builders.

The next *lab instrument* is a model.

Weather & climate

Monsoon onset, cloudbursts in the Himalaya, district-level forecasts. Models trained on physics plus decades of IMD and satellite data; India's own data is the asset nobody else has.

Materials & energy

Search millions of candidate compounds for batteries, solar cells and catalysts before touching a furnace. The power grid itself, with rooftop solar and storage, needs AI that respects physics: that is our lab's work.

Biology & medicine

Protein structure prediction changed drug discovery. India's disease burden, genetic diversity and hospital records are unstudied at scale. Whoever builds on them decides what the medicines of 2040 target.

Space & earth observation

ISRO's satellites produce terabytes a day. Crop yield, groundwater, glacier retreat, illegal mining: all readable from orbit if someone writes the model.

Mathematics itself

AI systems now solve olympiad-level problems and help check proofs. Your Class 11 maths is the entry ticket to a field that is being rebuilt in front of you.

Frugal science

A microscope on a phone. A sensor built for ₹500. AI that squeezes signal out of cheap hardware is an Indian research tradition, and a global export.

Science is becoming "data + physics + model". Bharat has the first two in abundance. The third is your job.

WHY IT COMES DOWN TO YOU

Three choices decide what an AI becomes: **what data** it reads, **what tests** it is judged by, and **which problems** its builders think are worth solving.

All three are made by people. If the people are in California, the model learns California. If they are from Roorkee, Ranchi and Rajkot, the model learns Bharat.

IF YOU WANT TO BUILD IT

The syllabus for building AI is *boringly old*.

- **Mathematics.** Probability (what you just saw), linear algebra (how a network stores patterns), calculus (how it learns). Your Class 11-12 maths is the foundation, not a hurdle.
- **Programming.** Python. Write code yourself for two years before letting AI write it for you.
- **A domain.** Power systems, biology, law, Sanskrit, soil science. AI is a tool; the problem comes from somewhere.
- **Language.** Your mother tongue is a technical asset now. English is still the language of the papers.
- **Writing.** Clear writing is clear thinking, and it is how ideas get funded.

Class 9-10

Maths basics,
first Python,
read a lot

Class 11-12

Probability,
matrices, one
real project

College

Pick a domain,
build for Bharat,
publish something

Work

Own the
judgement
tasks

LEAVE WITH SOMETHING TO DO

A *30-day* challenge.

Week 1

Attempt-first

Every homework: ten minutes alone before any AI. Keep a count of how many days you managed it.

Week 2

Catch a hallucination

Ask any chatbot for three references on a topic you know. Check each. Write down what it invented.

Week 3

Break the model

Ask the same question in English and in your mother tongue. Compare. That gap is a research problem.

Week 4

Build the toy

Write a 30-line Python program that does what this slide deck did: count word pairs, predict the next word. Send it to us.

REMEMBER ONE LINE

AI predicts the next *word*.
You decide the next *decade*.

Questions welcome!

Come find us at PSquare Lab, Electrical Engineering,
IIT Roorkee · psquare-lab.github.io