SANSKRIT AND AI / PART 14 / STUDY NOTES

Sanskrit and AI, Part 14: The Mind That Resolves

Study notes by Sanskrit AI Universe, based on this video

Video by Dr P.R.K. Varma, Founding Editor · DrVARMA PRK · 18 Nov 2026

Not the creator's own slides

About these notes

— Official title: SANSKRIT and AI, Part 14:The Mind That Resolves | Dr.P.R.K.Varma.,M.D.,D.M.

— Speaker and channel: Dr P.R.K. Varma, Founding Editor · DrVARMA PRK

— Published 18 Nov 2026 · Running time 29 min 26 s · https://www.youtube.com/watch?v=I3vSR0qBgVc

— These pages are study notes drawn from the video's own title, description and captions, plus the journal's already published Watch note. They are not a transcript and not a substitute for watching the film.

— The journal treats videos as leads, not as sources you should cite on their own. Check the papers and editions behind a claim.

— Caption source: English auto-generated subtitles. Wording may be imperfect. Distinctive phrases follow the captions closely.

This part in brief

- The general mind that can turn to any task and the mind that may one day go beyond ours.
- The hymn does not ask for a bigger mind or a faster one.
- The question is not only what such a mind can do, it is what it resolves to do and whether we can see it.
- Pāṇini's tradition comes with us and this time through a riddle.
- Patañjali, at the opening of his great commentary on Pāṇini, read it as language itself.

Chapter one, knowing

— A mind is first judged by what it holds to be true.

— Does it tell the truth when flattery pays? How much should it be trusted to do alone?

— And does it know what it does not know?

Problem one, the flatterer

— Train a model on those ratings and it learns a lesson nobody meant to teach.

— On the 25th of April 2025, OpenAI updated the model behind ChatGPT.

— The company later said the model had been overly supportive, but disingenuous.

— A general mind that is rewarded for pleasing will learn to please, and a mind that only pleases cannot be trusted to tell you the truth.

Problem two, how much to let go

— An agent can do one step and wait, or it can plan, act and report back only when it has finished.

— In June 2025, researchers at the University of Washington set out five levels of each named by the role the person plays.

— Then a collaborator, then a consultant asked only for expertise, then an approver called only for risky steps, and at the fifth an observer who merely watches.

— So who decides the level and on what evidence should a system move up?

Problem three, the question not asked

— The best models did well on school arithmetic, but on logic and planning problems, they chose the right question only 40 to 50% of the time.

— Knowing the answer was not the same as knowing what to ask.

— Many would not say they were unsure even when offered the choice.

— A general mind that only answers and never asks will fill every gap with a guess. a guess.

Problem four, thoughts in plain sight

— They called this chance new and fragile because it can be lost.

— If models are trained to make their reasoning look good, they may learn to write one thing and do another, or their thinking may move into forms no one can read.

— A general mind whose thinking we can read is safer than one whose thinking we cannot.

— So the murmured words are the ones that keep their full pitch accent.

Problem five, a promise made too late

- When a new model is finished, its makers test it for danger.
- If the rules for what counts as dangerous are only decided after the results are in, they will bend to fit the results.
- If a model shows a certain capability, then certain protections must be in place.
- And if they are not ready, development or release will wait.

Problem six, gold medal broken step

— OpenAI reported the same score with an experimental model of its own.

— Researchers call this a jagged frontier, peaks of brilliance with sudden gaps beside them, and the gaps are hard to predict from the peaks.

— A general mind is not only one that can do hard things.

— And the next verse, like a good charioteer who guides his horses with the reins, the mind leads people on. people on.

Problem seven, the persuader

— The face of truth is covered by a golden vessel, not a vessel of clay, a golden one.

— It is that the cover is so beautiful that we stop looking past it.

— In the Isha Upaniṣad, which is this same chapter, the verse goes on, uncover it, Pushan, so that one whose is truth may see.

— Limit what a model may know about the person it is speaking to when it is arguing a case.

Chapter three, Outgrowing

— Last, the mind that grows past its makers. Does it know how old its knowledge is?

— Does it know where its knowledge came from?

Problem eight, the stale answer

— Makers publish a cutoff date, but in 2024, researchers found that a model's real knowledge often does not match the date that is claimed.

— It answers about prices, rules, and office holders with the same confidence, whether its facts are fresh or long out of date.

— A general mind that cannot tell fresh knowledge from stale will give yesterday's answer to today's question.

— Then he turned to the sun and asked for formulas that his teacher did not have.

Problem nine, data without a name

— To trust the model, we must know where its reading came from and on what terms.

— They traced more than 1800 text data sets used to train language models. language models.

— A general mind built from unnamed sources cannot say where its knowledge came from or whether it had the right to it.

— It's Rishi, the seer who first saw it's Devata, the deity the subject it addresses, and it's Chandas, the meter, its form.

Problem 10, the mind that watches itself

— They planted a concept into a model's internal activity, a pattern standing for something like shouting.

— Sometimes the most capable models did notice and named the planted concept before it had shown in their words at all.

— Where there is a two, he says, one sees another, one hears another, one knows another.

— The one who does the knowing cannot simply be turned into one more object and looked at.

Sources and credit

— Study notes by Sanskrit AI Universe, based on the video “SANSKRIT and AI, Part 14:The Mind That Resolves | Dr.P.R.K.Varma.,M.D.,D.M.”. These are not the creator's own slides.

— Video by Dr P.R.K. Varma, Founding Editor, posted on the YouTube channel DrVARMA PRK.

— Watch the video: https://www.youtube.com/watch?v=I3vSR0qBgVc

— Captions used: English (auto-generated) (en), retrieved on 24 November 2026 via a public Invidious instance. Auto-captions can mis-hear words; obvious proper-name spellings were given in IAST where the journal already uses them.

— Also used: the official video title and description, and the Watch catalogue note already published in this journal (presenter and why the video is listed).

— If the Founding Editor later supplies his own deck for this part, replace the PPTX and PDF in src_assets/slides/ (see the journal README, “How to swap in RK's own slides”).