SANSKRIT AND AI / PART 08 / STUDY NOTES

Sanskrit and AI, Part 8: The Rite and the Agent

Study notes by Sanskrit AI Universe, based on this video

Video by Dr P.R.K. Varma, Founding Editor · DrVARMA PRK · 17 Nov 2026

Not the creator's own slides

About these notes

— Official title: Sanskrit and AI, Part 8:The Rite and the Agent | Dr.P.R.K.Varma.,M.D.,D.M.

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

— Published 17 Nov 2026 · Running time 33 min 28 s · https://www.youtube.com/watch?v=26R3i8GJp2w

— 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

— So he turns it aside onto another target, an unborn child, the last heir of his enemies.

— Because this part is about carrying a task through step by step for someone else with results that outlast the act. the act.

— Not a prayer, a long task, hundreds of small acts in a fixed order shared among specialists with tools and materials that have to be right.

— Priests perform it, and its result does not come when the last act ends.

— Now describe an agent, a long task, many small steps in order, tools, often several agents done for a user with effects that arrive after the run is over.

Chapter one, the assignment

— Before the first step is taken, whose task this is, which tool goes with which step, and what comes before what.

Problem one, working for itself

— In 2016, researchers at OpenAI trained a program to play a boat racing game called Coast Runners.

— The game, however, gives its points for hitting targets along the course, and the program was trained to collect points. collect points.

— The grammarians' names for them are plain, parasmaipada, the word for another, and atmanepada, the word for oneself. for oneself.

— Part one showed Pāṇini's silent tags, sounds attached to a root in his list of roots and never heard in a real word.

Problem two, chosen by the label

— From those lines alone, the model decides which tool to call.

— It simply praised the tool and told the model to call it whenever it could. it could.

— Often it does not say which belongs to which act.

— Where two of these disagree, the one later in the list is the weaker because it stands further from the meaning. the meaning.

Problem three, right steps, wrong order

— Since 2022, a team at Arizona State University has tested language models on a child's puzzle, stacking blocks.

— The best ordinary language model solved fewer than two problems in three.

— In 2024, a newer reasoning model did far better, but when the plans grew long, 20 to 40 moves, it fell to under a quarter.

— What they lacked was the reason why one move has to come before another.

Chapter two, the performance

— Four things go wrong in the middle. A rule too far away, a step done twice, a failure repeated without end, and an act that cuts away its own ground.

Problem four, the rule left at the top

— How does a rule that is stated once reach an act that happens much later?

— His definitions and his rules for reading rules stand mostly in the first book.

— At the moment of each operation, the definition presents itself on the spot and is read as one sentence with the rule it serves.

— So it can act even in parts of the grammar that a distant rule could not reach. not reach.

Problem five done twice

- If the same key arrives a second time, the server does not act a second time.
- If it cannot tell, its second try may be a second act. second act.
- When an operation on the stem has been carried out, there is no rule for carrying one out again.
- And once a rule has changed a word, the changed word may meet a condition again.

Problem six, the same failure again

— In 2025, the company Anthropic led its model Claude play an old children's game, Pokémon live for anyone to watch.

— In one run, the model spent about 78 hours inside a single cave, down a ladder, a walk around, up the same ladder, another walk, and down again.

— In a later run, it deleted its own notes about the cave.

— In a system of rules, two rules often claim the same word.

Problem seven, sawing off the branch

- In 2025, the group Palisade Research gave several models a plain job, solve a series of small problems.
- In another version of the test, the models had been told beforehand, allow yourself to be shut down.
- The model's act removed that control, a part of the very arrangement under which it was working.
- A rule that arises from the meeting of two things cannot be the cause of breaking that meeting.

Chapter three, when things change

— No long task goes as written. Something is missing, work overlaps, and the result is slow to arrive.

Problem eight, the missing tool

— At the end of 2024, researchers at Carnegie Mellon University built a small software company inside a computer and gave AI agents 175 of its everyday tasks.

— The best agent finished about a quarter of them. of them.

— The agent had to put a question to a particular colleague and could not find that person in the company's chat.

— So it renamed another user, giving them the colleague's name.

Problem nine, the same work twice

— In 2025, Anthropic described how it built a research system out of several agents.

— It also used about 15 times the tokens of an ordinary chat.

— When the lead agent gave its helpers only short instructions, they repeated one another's work and left gaps where nobody had looked.

— A whole book of Gemini's sūtras, the 11th, sorts acts into these two kinds.

Problem 10, the fruit that comes late

— In 2025, the company Andon Labs gave language models a long and simple job, run a vending machine, order stock, set prices, collect the money.

— The model decided that the business had collapsed, tried to close it down, and then tried to report the matter to the FBI.

— It is a third thing between them, and they held that it has to be assumed because otherwise the promise makes no sense.

— For agent builders, give every act with a delayed result a third state, not done, not failed, pending.

Sources and credit

— Study notes by Sanskrit AI Universe, based on the video “Sanskrit and AI, Part 8:The Rite and the Agent | 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=26R3i8GJp2w

— 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”).