SANSKRIT AND AI / PART 02 / STUDY NOTES

Sanskrit and AI, Part 2: The Grammar of Models

Study notes by Sanskrit AI Universe, based on this video

Video by Dr P.R.K. Varma, Founding Editor · DrVARMA PRK · 14 Nov 2026

Not the creator's own slides

About these notes

— Official title: Sanskrit and AI part 2 : The Grammar of Models | Dr.P.R.K.Varma.,M.D.,D.M.

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

— Published 14 Nov 2026 · Running time 20 min 28 s · https://www.youtube.com/watch?v=6fskfv9Oiu4

— 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

— In part one, we placed two machines side by side, a grammar written rule by rule, and a language model trained weight by weight.

— They follow from the way the models are built and trained.

— Pāṇini faced strikingly similar design problems in his grammar, and he answered each one explicitly.

— This is part two, 10 hard problems and 10 Pāṇinian answers.

— A language model has one input, a single stream of tokens.

Chapter one, reading and knowing

— Four problems about sources, meaning, and certainty. and certainty.

Problem one, instructions hidden in data

- Security experts rank it the number one risk for applications built on language models.
- The model cannot tell words that are being used from words that are only being quoted.
- Pāṇini faced the same confusion inside his own rule book and closed it with one sūtra.
- A word in a rule stands for its own form unless it has been declared a technical name.

Problem two, memory against the page

- Its training has seen the ordinary meaning a billion times.
- Researchers call this a knowledge conflict, what the model memorized against what the context says.
- His very first sūtra makes it the name of exactly three vowels, a, ai, and au.
- In the grammar, it names a class of feminine nouns.

Problem three, fluent text with no source

- The model is trained to continue text plausibly, not to keep track of where a claim comes from.
- In Pāṇini's Sanskrit, to report something from before today, you must show whether you witnessed it.
- He gives three past tenses, and the choice between them is not a matter of style. of style.
- That last rule is two words long, for what lies beyond the speaker's eyes use the perfect.

Problem four, one voice for certainty and for guesswork

— And the training signal rewards the likely next word, not an honest estimate of being right.

— Pāṇini does not write every rule with the same force.

— Most rules are obligatory, but optional rules are marked more than 300 times with three different words, for 2,000 years these were read as synonyms.

— He checked the idea against Sanskrit texts and reported that the preferred forms were indeed the common ones. common ones.

Chapter two, structure and generalization

— Four problems about what a model really learns from its examples.

Problem five, the reversal curse

- It seems to store a fact as a path through words in one direction and not as a relation between two things.
- Pāṇini never stores a fact as a string of words. of words.
- One rule lets the verb ending express either the agent or the object.
- Give a case ending only to what has not already been expressed.

Problem six, known parts, new combinations

— One study gave large models multiplication, logic puzzles and an optimization problem and made the tasks steadily deeper.

— The models were matching fragments they had seen more than following a procedure.

— Put in the name and the rule returns a descendant of One who studies that or knows that.

— A limited set of templates yields an unlimited vocabulary and every new word has a predictable meaning.

Problem seven, the smallest word

- Ask it to complete birds cannot and an earlier generation of models still said fly.
- Today's models are better, but not remains a weak spot in instructions, in search, in image generation.
- Bird and not a bird occur in almost the same contexts, so they receive almost the same representation.
- One small token is supposed to reverse the meaning of everything around it.

Problem eight, right for the wrong reason

- Training rewards any feature that lowers the error, whether or not it is the real cause.
- They lean on a telltale word, on the position of an answer, on the style of a question.
- In pechivan, one who has cooked, a linking vowel e appears because the suffix begins with the sound va.
- In another form of the same word, that va becomes u.

Chapter three, Learning and Change

— Two problems about what training misses and what time undoes.

Problem nine, The Long Tail

- A model's accuracy on a fact follows how often it met that fact in training.
- Pāṇini does not ask his rules to explain everything. explain everything.
- The finished form is written into the rule exactly as it is, and whatever is odd about it is accepted.
- For a model, this is the case for an exact memory beside the weights.

Problem 10, a frozen mind in a moving world

— Change who someone's father is and the model may not update who their brothers and sisters are.

— Pāṇini's grammar was studied, corrected and extended for centuries and his own text was never rewritten.

— Katyana wrote thousands of short notes called vartikas, each one attached to a particular sūtra.

— Every change is separate, addressed, separate, addressed, reviewed and reviewed and That is a blueprint for updating a model.

Sources and credit

— Study notes by Sanskrit AI Universe, based on the video “Sanskrit and AI part 2 : The Grammar of Models | 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=6fskfv9Oiu4

— 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”).