



Foundations First: How Will Mathematics and Statistics Shape India's AI Ambitions?*

Sudeshna Guha Neogi¹, Abhijit Kulkarni², Ramasubramanian Sundararajan³,
Balaji Raman⁴ and Sourish Das⁵

¹SAS, ²Principal Financial Group, ³SOLUS.ai,
⁴Cogitaas, ⁵Chennai Mathematical Institute

Received: 03 September 2026; Revised: 05 October 2026; Accepted: 05 October 2026

Abstract

Artificial Intelligence has become the organising ambition of Indian higher education and industry alike. Yet the enthusiasm filling our AI and data science classrooms is thinning the foundational training in probability, statistical inference and mathematics on which those systems rest. This note argues that we mistake fluency with tools for competence in reasoning; that the properties we now demand of AI systems, namely explainability, validation and safety, are statistical properties whatever language we use for them; and that India's national AI ambitions fund the application layer far more generously than the foundations beneath it. We state plainly what we are not arguing, including where machine learning practice has outrun statistics itself. We close with six implementable recommendations, each naming the body that would have to act.

Key words: Statistics education; Artificial intelligence; Curriculum design; NEP 2020; IndiaAI Mission; Statistical consulting.

AMS Subject Classifications: 62K05, 05B05

*This is a short communication and set of recommendations arising from a panel discussion at the Annual Conference of the Society of Statistics, Computer and Applications, Pune, 2026

1. Introduction

When we convened this panel, we set ourselves a deliberately uncomfortable question: if AI is the future, why does the discipline that underwrites it feel like it is losing ground?

The worry is easy to state. Across our universities, students move in large numbers toward applied AI and machine learning courses, while curricula increasingly treat probability theory, statistical inference, linear algebra and optimisation as obstacles to clear rather than foundations to build on. Students respond rationally to what the job market appears to

reward. In the panel’s collective experience they also make a slow-acting mistake, one whose costs will not appear in this year’s placement statistics but in the next decade’s capacity to do original work.

We should be candid at the outset about the standing of that claim. This note reports the considered experience of practitioners and teachers rather than the results of a systematic study. We return to this limitation, and to what would settle the question, in Section 7.

2. What does AI reliability actually ask for?

The central observation draws a distinction between fluency and understanding. A great many graduates today can do remarkable things: assemble a pipeline, fine-tune a pre-trained model, deploy a working system in an afternoon. These are real skills and it would be foolish to disparage them. But fluency with tools is not the ability to reason about a problem, and the gap between the two is precisely where statistics lives.

As foundations thin, students lose a particular set of intellectual moves: posing a problem as an estimand, so that they know which quantity they are trying to learn before they reach for a method; asking whether the data at hand identify that quantity at all; quantifying uncertainty honestly rather than reporting a point estimate as though it were the truth; distinguishing correlation from causal structure; recognising misspecification and knowing what it does to their conclusions. None of these is exotic. They constitute the ordinary discipline of statistical thinking (Cobb, 2015; Donoho, 2017). But they stay invisible in a workflow that treats held-out accuracy as the only criterion, and they erode when departments demote conceptual courses to optional status. The profession has learned this lesson painfully once already: the American Statistical Association issued a formal statement precisely because practitioners across the sciences had reduced inference to a single number they no longer understood (Wasserstein and Lazar, 2016).

Meng (2018) makes the sharpest available case that this is not a matter of taste. His analysis of the big data paradox shows that a very large but slightly non-representative sample can carry less information than a much smaller random one, and that anyone who reasons about sample size rather than data quality will miss the error entirely. That is a foundational insight, and no amount of tooling substitutes for it.

The strongest objection deserves a hearing. One might argue that a practitioner can always acquire foundations later, when the work demands them, and that meanwhile tool fluency creates genuine value today. There is something to this. But it underestimates how hard it is to retrofit conceptual depth onto someone trained to skip it, and how rarely working life creates the space to do so. Foundations are not a module one bolts on. They form a habit of mind, most cheaply acquired early. A practitioner who has never internalised the question “what am I actually estimating, and with what uncertainty?” is unlikely to start asking it under deadline pressure.

3. What are we not arguing?

The deeper point is that the qualities we now demand of AI systems, that they be explainable, validated and safe, are statistical qualities whatever language we dress them in. The panel prefers to demonstrate this rather than assert it, so we name the machinery.

Explainability. Practitioners often present explainability as a layer bolted onto a finished model. In substance it is statistical reasoning: sensitivity of outputs to inputs, identification of high-influence observations, the behaviour of one variable when others stay fixed, and whether the model’s assumptions hold at all. Rudin (2019) argues further that for high-stakes decisions the post-hoc explanation of an opaque model is often the wrong goal, and that an interpretable model built for the purpose serves better and loses little or no accuracy. Anyone without a modelling vocabulary cannot make that argument, or evaluate it.

Validation. Validation tells the same story about out-of-sample behaviour. Cross-validation, calibration of predicted probabilities against observed frequencies, and evaluation under strictly proper scoring rules (Gneiting and Raftery, 2007) are not bureaucratic formalities. They carry the scientific content of the claim that a system works. Conformal prediction (Angelopoulos and Bates, 2021) now supplies distribution-free coverage guarantees for the predictions of arbitrary black-box models, including large language models, and stands among the clearest recent demonstrations that a classical statistical idea can answer a modern engineering problem.

Safety. Safety asks about tails. As systems shift toward large generative models, the properties that matter most, namely rare but harmful outputs, behaviour under distribution shift, and drift over time, are properties of probability distributions and of risk. Nobody can construct a serious guardrail without thinking distributionally, which is to say statistically.

Reproducibility and auditability belong on the same list. Both demand a documented chain running from data to estimand to method to conclusion, and both break down when nobody in the room can state what the model estimates.

The erosion of foundations is therefore not a parochial concern of statistics departments protecting territory. It bears directly on whether the next generation can make AI trustworthy, which is the only kind of AI worth deploying in healthcare, finance, public policy and the other domains that operate under heavy regulation.

4. What we are not arguing?

The previous section presents an argument by statisticians that statistics should be compulsory, so we should declare the interest and mark its limits. Three clarifications follow.

We are not defending the inherited syllabus. The case for foundations is not a case for preserving a 1970s statistics curriculum in amber. Many Indian programmes still teach compulsory statistical content as a catalogue of named tests and closed-form distributions, and a graduate who has memorised that catalogue is not thereby prepared for anything. The foundations we mean are probability, linear algebra, optimisation, computation, calibration and causal reasoning, which teachers should deliver with data and code in the room (Nolan and Temple Lang, 2010; De Veaux *et al.*, 2017). Where the classical syllabus obstructs that, departments should cut it rather than defend it.

We do not claim that statistics holds all the answers. Modern machine learning has repeatedly outrun statistical theory. Heavily overparameterised models generalise well in regimes where classical bias-variance reasoning predicts they should not, and the

double descent phenomenon (Belkin *et al.*, 2019) now occupies statisticians precisely because practice found it first. Breiman (2001) warned the profession about this asymmetry a quarter of a century ago and events have proved him substantially right. The traffic between the two fields runs in both directions. Students will not believe a curriculum that casts statistics as the adult supervising an unruly engineering discipline, and they should not.

We are not arguing against AI coursework. Nothing in this note asks for fewer AI and machine learning courses. It asks that they rest on strong foundations.

5. The Indian inflection

This argument carries particular resonance here. India is the country of the Indian Statistical Institute and of a statistical tradition, running through P. C. Mahalanobis, C. R. Rao, R. C. Bose and those who followed, that built large-scale statistical practice into the machinery of national planning long before “data science” was a phrase (Mahalanobis, 1965; Ghosh *et al.*, 1999). It would be a strange irony if the nation that gave the world some of its foundational statistical thinking now produced graduates fluent in calling an API but unsure how to define an estimand.

The present moment opens a door as much as it poses a threat. The National Education Policy 2020 (Ministry of Education, 2020), with its emphasis on multidisciplinary and flexible undergraduate programmes, admits two readings: as a licence to let foundational courses wither into electives, or as room to insist that every student of AI also learns to reason about uncertainty. Boards of studies will settle which reading prevails, course by course, over the next few years.

The national AI programme sharpens the point. The Union Cabinet approved the IndiaAI Mission in March 2024 with an outlay of Rs 10,371.92 crore across seven pillars, covering compute capacity, datasets, applications, startups, safe and trusted AI, and skilling (Press Information Bureau, 2024). Its FutureSkills pillar expands AI course offerings at undergraduate, postgraduate and doctoral level, establishes Data and AI labs in Tier 2 and Tier 3 cities, and routes doctoral fellowships through the top 50 NIRF-ranked institutes (IndiaAI, n.d.). This is a serious and welcome commitment. Our observation is narrow and, we think, uncontroversial: the mission funds the AI layer generously and the mathematical foundations beneath it only incidentally, and no pillar takes responsibility for strengthening doctoral training in statistics and the mathematical sciences.

We would put the underlying asymmetry this way. Compute is capital: expensive, constrained, but purchasable, and India is buying it. No budget line buys the capacity to interrogate a method, recognise when it fails and invent a better one. That capacity grows only from sustained engagement with probability, mathematics and inference, over a decade, in departments that have the people to teach it. It marks the difference between a country that uses AI and one that advances it.

6. Recommendations

The panel held that diagnosis without prescription is cheap. What follows stays deliberately small and implementable, and each recommendation names the body that would have to act.

R1. Make foundations prerequisites, not electives. *For university boards of studies, with AICTE, UGC and NBA.* Every AI and data science degree should require a defined floor of foundational coursework, which students complete before their first deployment-heavy elective. We suggest a minimum of twelve credits spanning probability, statistical inference, linear algebra and optimisation. A board of studies can adopt or contest a number. It can do neither with “adequate grounding”.

R2. Institutionalise applied statistical reasoning through a consulting course.

For departments. A final-year statistical consulting or collaborative problem-solving course of three to four credits, which draws on live problems and pairs an academic supervisor with a practitioner. Such a course teaches what lectures cannot: how to turn a messy question into a formal one, how to communicate uncertainty to a non-specialist, how to defend a modelling choice, and how to take responsibility for a recommendation. The panel regarded this as the single highest-return reform available to us.

R3. Reform assessment, not only syllabus. *For examination boards and course instructors.* Prerequisites examined by rote produce gatekeeping without understanding. At least one third of the assessment weight in foundational courses should rest on open-ended items: formulate the estimand for a described situation, critique a flawed analysis, defend a modelling choice orally.

R4. Build a structured two-way relationship with industry. *For departments and industry advisory boards.* Advisory input into curricula, practitioners co-teaching, faculty engaging with real industrial problems, and industry problems feeding the course in R2. One guardrail: academic bodies retain final authority over curriculum content, so that departments achieve relevance without surrendering the conceptual core to short-term market signals.

R5. Fund the people who carry the foundations. *For IndiaAI FutureSkills, DST, UGC and universities.* Doctoral fellowships in statistics and the mathematical sciences, not only in AI; faculty development programmes for teachers of foundational courses; and recognition for faculty who genuinely bridge theory and practice. Teaching capacity, not policy intent, currently binds R1. Mandating prerequisites without it will produce failure rates rather than understanding.

R6. Carry this discussion into a standing, multi-institution process. *For professional societies, academic associations and curriculum bodies in the statistical, mathematical and computational sciences.* No single panel can settle a question of this scope. These recommendations should be taken up through further panel discussions at conferences and on other academic and industry platforms, and consolidated by a standing curriculum committee drawn from several institutions. Such a process could produce a model syllabus and reading list for the consulting course in R2, faculty development workshops attached to annual conferences, a shared repository of consulting case problems, and a student consulting competition. The audit described in Section 7 would give it an evidence base.

7. Limitations and an invitation

This note reports one panel at one conference. Its central empirical claims, that foundational coursework is losing ground and that the loss carries consequences, reflect the collective professional experience of the panellists, and this note does not establish them with systematic evidence.

They are, however, testable. A curriculum audit of twenty to thirty Indian AI and data science programmes, recording compulsory credits in probability, inference, linear algebra and optimisation together with their position in the degree sequence, would settle the descriptive question within a few months. Programme documents are largely public, the All India Survey on Higher Education (Ministry of Education, 2024) supplies the enrolment denominator, and the AICTE model syllabi (AICTE, n.d.) supply a benchmark. We invite researchers across the statistical and computational sciences to undertake that audit, and we offer this note as the hypothesis it would test. A workable protocol would be the following.

1. **Frame.** Draw a stratified sample of twenty to thirty programmes spanning IITs and NITs, central universities, state universities and private universities, and covering B.Tech in computer science with an AI or ML specialisation, B.Tech in AI and data science, and MSc in data science or statistics.
2. **Record.** For each programme, record total credits, compulsory credits in probability, statistical inference, linear algebra and optimisation, and the semester in which each sits relative to the first deployment-heavy course.
3. **Classify.** Count a course as foundational only when its stated objectives include derivation or proof, so that tool-oriented service courses do not inflate the count.
4. **Benchmark.** Compare against the AICTE model curriculum and against De Veaux *et al.* (2017).
5. **Test displacement, not level.** Where programme documents survive from an earlier year, repeat the exercise at two time points. The hypothesis concerns change over time, and a single snapshot cannot refute it.
6. **Report.** Publish the coding rules and the programme list, report proportions with uncertainty, and describe the sample as a convenience sample unless a complete frame is used.

8. Beyond the hype cycle

AI moves in cycles of enthusiasm; foundations endure across them. Probability, statistics and mathematics have outlasted many waves of fashionable technique and will outlast this one, because they are not techniques but ways of reasoning under uncertainty. The task before this community is to hold institutional nerve: to embrace AI wholeheartedly while refusing to let its excitement hollow out the disciplines that make it reliable.

How well we reason, and not only how much we compute, will decide the future of AI in India. That is a statistical question, and answering it well is our responsibility.

Acknowledgements

We thank the organisers of the SSCA Annual Conference, Pune, 2026 for convening the panel, and the participants from the floor whose questions shaped several of the arguments above. The authors are grateful to the anonymous reviewer whose suggestions have helped in improving the presentation of the manuscript.

Conflict of interest

The authors do not have any financial or non-financial conflict of interest to declare for the research work included in this article. All authors are statisticians and hold professional and institutional interests in the standing of statistics within Indian higher education. We believe the argument stands on its merits, and we record the interest so that readers can judge for themselves.

References

- Angelopoulos, A. N. and Bates, S. (2021). A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv preprint* arXiv:2107.07511. <https://arxiv.org/abs/2107.07511>
- All India Council for Technical Education (n.d.). *Model Curriculum for Undergraduate, Postgraduate, Diploma and Minor Degree Programmes in Engineering and Technology*. <https://www.aicte-india.org/education/model-syllabus>
- Belkin, M., Hsu, D., Ma, S. and Mandal, S. (2019). Reconciling modern machine-learning practice and the classical bias-variance trade-off. *Proceedings of the National Academy of Sciences*, **116**, 15849–15854. <https://doi.org/10.1073/pnas.1903070116>
- Breiman, L. (2001). Statistical modeling: the two cultures (with comments and a rejoinder by the author). *Statistical Science*, **16**, 199–231. <https://doi.org/10.1214/ss/1009213726>
- Cobb, G. W. (2015). Mere renovation is too little too late: we need to rethink our undergraduate curriculum from the ground up. *The American Statistician*, **69**, 266–282. <https://doi.org/10.1080/00031305.2015.1093029>
- De Veaux, R. D., Agarwal, M., Averett, M. *et al.* (2017). Curriculum guidelines for undergraduate programs in data science. *Annual Review of Statistics and Its Application*, **4**, 15–30. <https://doi.org/10.1146/annurev-statistics-060116-053930>
- Donoho, D. (2017). 50 years of data science. *Journal of Computational and Graphical Statistics*, **26**, 745–766. <https://doi.org/10.1080/10618600.2017.1384734>
- Ghosh, J. K., Maiti, P., Rao, T. J. and Sinha, B. K. (1999). Evolution of statistics in India. *International Statistical Review*, **67**, 13–34. <https://doi.org/10.1111/j.1751-5823.1999.tb00378.x>
- Gneiting, T. and Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, **102**, 359–378. <https://doi.org/10.1198/016214506000001437>
- IndiaAI (n.d.). *IndiaAI FutureSkills*. Ministry of Electronics and Information Technology, Government of India. <https://indiaai.gov.in/hub/indiaai-futureskills>

- Mahalanobis, P. C. (1965). Statistics as a key technology. *The American Statistician*, **19**, 43-46. <https://doi.org/10.1080/00031305.1965.10479715>
- Meng, X.-L. (2018). Statistical paradises and paradoxes in big data (I): law of large populations, big data paradox, and the 2016 US presidential election. *The Annals of Applied Statistics*, **12**, 685–726. <https://doi.org/10.1214/18-A0AS1161SF>
- Ministry of Education, Government of India (2020). *National Education Policy 2020*. New Delhi: Government of India.
- Ministry of Education, Government of India (2024). *All India Survey on Higher Education (AISHE) 2021–22*. Press Information Bureau release, 25 January 2024. <https://www.pib.gov.in/PressReleasePage.aspx?PRID=1999713>
- Nolan, D. and Temple Lang, D. (2010). Computing in the statistics curricula. *The American Statistician*, **64**, 97–107. <https://doi.org/10.1198/tast.2010.09132>
- Press Information Bureau, Government of India (2024). *Cabinet Approves Ambitious IndiaAI Mission to Strengthen the AI Innovation Ecosystem*. 7 March 2024. <https://www.pib.gov.in/PressReleaseIframePage.aspx?PRID=2012355>
- Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, **1**, 206–215. <https://doi.org/10.1038/s42256-019-0048-x>
- Wasserstein, R. L. and Lazar, N. A. (2016). The ASA statement on p-values: context, process, and purpose. *The American Statistician*, **70**, 129–133. <https://doi.org/10.1080/00031305.2016.1154108>