THE INDIA FORUM
A JOURNAL-MAGAZINE ON CONTEMPORARY ISSUES

November 20, 2026

India's Computing Boom Is Passing Its Scientists By

By: Sameer Saurav

India is spending heavily on computing power, and much of today's science runs on computers. Yet the money is organised around AI rather than the sciences that use it, and access is easiest for researchers at large, well-funded institutions.

As a chemistry PhD student at IIT Kanpur, I rarely gave a thought to where my computing time came from. Modern research in the basic sciences increasingly relies on supercomputers that can process up to thousands of trillions of calculations in seconds and it was natural for me at a premier institute to take this aid for granted.

> Computing is now one of the ordinary ways of doing research, and it could let good science happen well beyond the campuses...

Some of my calculations ran for weeks on hundreds of processor cores. I wrote a job script, sent it to the institute's cluster of linked processors (in effect, its supercomputer), and came back to it a few days later. When a job crashed, someone down the corridor usually knew why and could fix the problem.

I now teach chemistry at a college in a small town in Bihar. My training has not changed, but I no longer have a cluster to continue working on my research questions. I am not alone. Many people with a PhD from India's top research institutions take teaching jobs outside them, and if they want to keep doing computational research, they have to find the computing time themselves.

If computing were a small corner of science, this would be a minor problem. But computing is now one of the ordinary ways of doing research, and it could let good science happen well beyond the campuses where it is concentrated today. Whether India gains from that will depend less on how many machines it buys than on who is allowed to use them.

A different way of doing science

For most of its history, chemistry was done with glassware. A chemist who wanted to know how fast a reaction ran would mix the reactants, measure how the mixture changed, and work out from that what the molecules must be doing.

Computer simulations replaced the flask with a model. The chemist describes the molecules with the laws of physics and the computer solves the equations. Depending on the question, it may work out the energy and shape of a single molecule, or follow thousands of atoms as they move and react. From the results one can read off energies, structures, rates, and other properties.

Experiments can often measure these quantities too, but they rarely show how a change actually happens, step by step. A simulation can show that mechanism. The same approach, with different equations, is now used across science, from biology to climate.

None of this makes the laboratory obsolete, because simulations still have to be checked against experiments. What has changed is how the laboratory is used. Drug designers now screen large libraries of candidate molecules on a computer, and chemists make only the few dozen that look promising. Materials scientists work out which compounds might make a better battery or catalyst before they synthesise any. Climate scientists cannot experiment on the planet, so for them the model is the laboratory. Biologists use simulations to watch proteins fold and enzymes open and close, motions too small and too fast for a microscope. The Nobel Prize in Chemistry went to computer models of complex chemical systems in 2013, and in 2024 to computational protein design and to AlphaFold, which predicts protein structures. The 2024 prize also shows how close simulation and AI have come.

> Good computational work can be done from a small-town college as well as from a big institute, as long as the researcher can get time on a large computer.

My own PhD asked how fast a drug molecule leaves its binding site, the pocket on its target where it sits and does its work. This matters because how long a drug stays bound can decide how well it works. A generation ago, that rate could only come from an experiment, and a computer could offer little more than a static picture. The difficulty is time. A simulation that follows moving atoms

advances by about a femtosecond, a millionth of a billionth of a second, at each step, while leaving the pocket can take a millisecond or longer. That means around a trillion steps. My work was on methods that get around this, so that such rates can be computed in weeks rather than ages.

In practice the work goes round in a loop. I build the model, with the drug, its target, and several thousand water molecules, and test it. Then I run dozens of simulations, each taking days on its own set of processors. When I analyse the results, I often find that something is off, and I go back and change it. So what I need is not one enormous job but hundreds of medium-sized ones spread over a year. They come in bursts, and each round depends on what the last one showed. Much of computational chemistry and biology works this way.

Computational science is also cheap to set up compared with most research. A wet laboratory can cost crores of rupees and take years to build. A computational group needs a trained person, software that is often free, and computing time. So, in principle, good computational work can be done from a small-town college as well as from a big institute, as long as the researcher can get time on a large computer. Getting that time is the hard part.

What India Lacks

Most of India's public computing comes from two programmes. The National Supercomputing Mission (NSM) was launched in 2015 with an outlay of 4,500 crore rupees and is carried out by C-DAC and the Indian Institute of Science (IISc) at Bangalore. By late 2025 it had installed 37 supercomputers, adding up to about 40 petaflops (DD News 2025), mostly at IITs, IISERs, and national institutes. (A petaflop is a thousand trillion calculations a second. If every one of India's roughly 145 crore people did one sum a second, it would take them about eight days to do what a one-petaflop machine does in a single second.)

The host institutions hand out time, mostly to their own researchers and at little cost. The government says the mission supports more than 10,000 researchers from over 200 institutions, who have run more than a crore of computing jobs and published over 1,500 papers, and that it has brought supercomputing to researchers in Tier II and Tier III cities. That is a real achievement, even though the mission has fallen short of its original plan for more than 70 machines. Most of the machines run at over 85% of capacity and many above 95%.

By world standards, 40 petaflops is modest. On the June 2026 TOP500 list, China's LineShine, the fastest machine in the world, came in at about 2,200 petaflops, followed by El Capitan in the United States at about 1,800. Germany's JUPITER has reached 1,000.

> The IndiaAI scheme subsidises a market built for AI, and who gets in depends on eligibility rules, paperwork, and whether one can pay the rest of the bill.

India's fastest public research machine, AIRAWAT at C-DAC Pune, is 211 on the list. It entered the list at 75th in 2023 and has slipped on every list since, as other countries have added faster machines. All the NSM's machines together come to less than a fortieth of El Capitan. But the NSM was never meant to build one giant machine. It set out to build many mid-sized ones, which suits most academic work. The problem, rather is that those machines run close to full capacity, and nobody publishes how long outside users wait, how many requests are refused or how much time goes to people outside the host institutions. Without those numbers, it is hard for anyone outside the system to say whether capacity has held research back.

The second programme is the compute arm of the Rs 10,372-crore IndiaAI Mission which works differently. It measures capacity differently in numbers of graphics processing units (GPUs), processors first built to draw video-game graphics. A GPU can do thousands of simple calculations at once, which makes it well suited to training AI models and to many kinds of scientific simulation.

Unlike in the NSM, the government does not own any GPUs itself. Instead, it pays part of the bill when approved users rent them from private cloud firms. More than 38,000 GPUs are now on its portal, and eligible users can get up to 40% of their bill paid.

Officials describe this as democratising computing, and the list of people who can apply, from startups to teaching faculty, is certainly long. The NSM, in other words, is public infrastructure built for science, with time given out by the institutions that host the machines. The IndiaAI scheme subsidises a market built for AI, and who gets in depends on eligibility rules, paperwork, and whether one can pay the rest of the bill.

How Much Computing is That?

GPU counts are misleading in a certain way. Many scientific calculations, quantum chemistry among them, need "double precision", numbers carrying about sixteen digits, because tiny rounding errors pile up over billions of steps. AI applications, on the other hand, can manage with much rougher numbers. This means a GPU's headline speed overstates what it can do for much of scientific work. NVIDIA's H100, one of the most powerful GPUs in the world, can do only 34 to 67 teraflops in double precision but nearly 1,000 teraflops in the 16-bit format used for AI.

Even if 10,000 of the 38,000 GPUs on the IndiaAI portal were high-end chips like the H100, in double precision, those alone would add up on paper to 340 to 670 petaflops, or roughly 8 to 17 times the whole NSM. In reality, the portal has a mix of models, and some, such as NVIDIA's L4 and L40S, are built mainly for running AI models and have little double-precision ability. These are peak estimates, and the GPUs are spread across several providers rather than wired into one machine.

Still, the conclusion holds. Even when we count only double-precision work, the kind most useful to science, the GPUs the government is paying for under the AI Mission add up to many times more computing than all the supercomputers it has built for science under the NSM. India's fastest machine on the June 2026 list of supercomputers, after all, is not a public one. It is Shakti Cloud, run by Yotta Data Services, one of the private firms that rent GPUs to the IndiaAI Mission, at 32nd place with about 84 petaflops (TOP500 2026). On its own, it is roughly twice the size of the entire NSM.

Where the AI compute went

The IndiaAI Compute Portal publishes its approvals, and in October 2026, *The Print* went through all 304 of them. Of about Rs 1,275 crore rupees of subsidies, 83% went to 16 proposals for foundation-model builders-teams building large, general-purpose AI models of the kind behind chatbots such as ChatGPT. In these cases, the government paid nearly all of their bills (98.9%), an exception to the usual 40% limit made for teams it had chosen to build Indian AI models. Most of that funding went into building Indian-language models at the IIT Bombay-led BharatGen, Sarvam AI, Gnani.ai, Soket AI, amongst others. Government departments and public sector undertakings took 12.5% of the subsidies, and fifteen government bodies, including the Rajya Sabha and the CBI, had about Rs 39 crore worth of computing time paid for in full.

Researchers, students and fellows received about Rs 10 crore, 0.8% of the total, and paid around 60% of their own bills. Even within the "researchers or academia" category, about Rs 163 crore of the total of Rs 173 crore went to a single approval for a startup. Foundation-model grants were approved, on median, four days before their start dates. Researcher grants were approved 29 days after theirs.

In the public reporting I have seen, I could not find a large allocation for chemistry, materials, climate, or biology. "AI" can mean almost anything, but here it has mostly meant language models and the digital services of government.

It is not the case that there is high utilisation of the GPUs being used, as in the case of the NSM machines. On the contrary, *The Print* cites a *Business Standard* report that the government has considered letting suppliers sell capacity abroad because there is "not enough uptake" in India.

> One of the things the committee is asked to weigh is a project's "national importance", measured by whether it may reach ten lakh citizens or have an economic impact of Rs 500 crore or more.

So there is spare capacity, and there are scientists who could use it, but they do not fit the scheme. The reason lies in how the scheme is designed. The researcher category is written for AI researchers and eligibility depends on an h-index of at least 5 (or on papers in AI or "allied" journals with an impact factor above 2.5).

Whether by design or by default, the scheme favours applicants who are large, well funded, and already working on AI. This is reinforced by a provision where startups can be asked to show revenue or investor funding. One founder who wanted compute worth under Rs 10,000 rupees was asked, at two days' notice, to prove revenue of at least Rs 50 lakh rupees or funding of at least Rs 1 crore, *The Print* reported.

It is also unclear whether the IndiaAI Mission has the expertise to evaluate a proposal in computational chemistry. Large requests (above 5,000 GPU hours or 50 GPUs) and all requests for a subsidy go to a Program Monitoring and Evaluation Committee, but the policy documents do not say who sits on it. (The NSM, on the other hand is overseen by a national steering committee and a technical advisory committee (IMPRI 2025), and time on its machines is allocated by the host institutions.)

More tellingly, one of the things the IndiaAI Mission committee is asked to weigh is a project's "national importance", measured by whether it may reach ten lakh citizens or have an economic impact of Rs 500 crore or more. Basic research rarely passes a test like that.

Democratisation or Rationing?

For an individual scientist, the problems with applying for computing time from the IndiaAI Mission are practical ones. When the portal was launched, the IT minister put the price at around Rs 67 per GPU hour after the 40% subsidy. At that rate, 3,000 GPU hours would still leave a college teacher paying about 2 lakh rupees, with no budget head to pay it from. One could reply that the teacher should simply apply for a research grant, for instance from the Anusandhan National Research Foundation (ANRF). That is possible, but a grant is competitive and takes months to come through, and a strong proposal usually needs preliminary results, which themselves need computing time.

The ANRF is also still settling in. A parliamentary committee noted in March 2026 that it had spent nothing in its first two years, and only about three-fifths of its 2025-26 budget so far. Even with a grant, the money may not fit. Project budgets are split into heads such as equipment, manpower, and consumables, and an IIT faculty member told *The Print* that cloud computing can only be bought as a consumable, an amount that is "restricted by design". A grant is the right route for a full project, but it is a slow and uncertain way to get the first few thousand hours needed to test an idea. Even at an IIT, a faculty member told *The Print*, registration "takes more than a month".

> The "democratisation" of computing through the IndiaAI mission has meant allowing more people to apply but then letting cost, paperwork, and timing to decide who actually gets compute.

Approvals also lapse if the time is not used within 30 days. That suits a startup training a model to a deadline, not the slow loop of simulation. Software is another hurdle. On a university cluster, staff install the common chemistry and physics packages. On rented time, the researcher has to do it alone, sometimes with commercial licences, and pays for every hour spent getting it to work. Omprakash Subbarao, who heads FSID CORE, IISc's research and development arm, has noted that a first-come-first-served system favours institutions that already have people who are good at writing proposals.

Building Indian-language models is a reasonable national goal. But the IndiaAI mission treats computing as an input to AI products, and it measures success by models launched and startups served. Science comes in only as a possible user of AI. A policy built this way will tend to favour quick, visible results over research that may take a decade to pay off.

In practice, the "democratisation" of computing through the IndiaAI mission has meant allowing more people to apply but then letting cost, paperwork, and timing to decide who actually gets compute. Not surprisingly, the money has gone to those best placed to apply. Some of those projects may well be worthwhile, but the scheme rewards the ability to apply and to pay, not the quality of the research question. On paper, it is open to everyone. In practice, it is not.

> A market, like the IndiaAI mission is, can work when everyone is able to pay, but college researchers in India are not. For

Rationing, the NSM's approach, suits research that cannot pay upfront, but it has problems of its own. The ration is small for a country of India's size. Host institutions look after their own researchers first, and outsiders have to apply individually to each computing centre. Rules for outside users differ from centre to centre. Some centres charge outsiders, like at the IUAC Delhi, which was set up by the UGC to serve universities. Outsiders who run into problems find little help, and the mission has been slow to add new machines.

A market, like the IndiaAI mission is, can work when everyone is able to pay, but college researchers in India are not. For them, rationing is the better model, as long as the rules are public and the same for everyone.

Experiences Abroad

Other countries have made rationing work by giving researchers clear entitlements. In the United States, the National Science Foundation's ACCESS programme gives small allocations on the basis of a short research plan, sometimes within a couple of working

days. In the Netherlands, SURF does the same within one to two weeks. Large requests go to peer-review by scientists in the field.

> China spends on a scale India cannot match, keeps much of its capacity off public lists, and gives its vouchers mostly to companies.

China runs national supercomputing centres in cities such as Tianjin, Guangzhou, Shenzhen, and Wuxi as public facilities for universities, government, and industry. Since 2023 it has been linking them through a national network, SCNet, which was unveiled in April 2024. By late 2025, SCNet connected more than 30 computing centres, served over 1,000 institutions, and offered nearly 100 ready-made scientific workflows. A national policy in December 2024 encouraged cities to hand out "computing vouchers". Shanghai covers up to 80% of AI computing rental for small firms, and Chengdu has set aside 100 million yuan for research institutions. China has also kept building machines for traditional science. LineShine, now the fastest machine in the world, uses no GPUs at all, and it also tops a benchmark that comes closer to how real scientific programs behave.

A comparison with China has clear limits. China spends on a scale India cannot match, keeps much of its capacity off public lists, and gives its vouchers mostly to companies. Still, India can learn from it in two ways. China treats computing for science as public infrastructure, separate from its AI policy, and it is bringing its scattered machines under one system that researchers can reach in one place. India has scattered machines too, but it has done neither of these things.

What is At Stake

The first cost is wasted training. India spends a lot to train computational scientists, and many of them-like me-join colleges and universities away from the major research institutions. Teaching rightly comes first in these jobs, and research has to fit around it. A determined person can manage that, but cannot work around the lack of a powerful computer. Without one, years of specialised training slowly go to waste.

The second cost is to innovation. New ideas tend to come from many people trying many things. In theory, researchers trained at IITs, IISERs, and national laboratories moving to district towns, along with cheaper computing, should spread Indian science more widely. They could also bring local problems within reach of the people who live with them.

> A teacher with a PhD and a good research question, working far from the big institutes, still has no straightforward way to get the computing time needed to run a calculation.

Screening molecules for diseases common in India, designing cheaper materials for batteries and solar cells, and modelling how a changing climate will affect a particular region can all be done with a computer from anywhere. This has not happened, because the researchers moved but access to computing stayed where it was. Their students also miss the chance to see current research being done up close, and fewer of them think of doing it later.

The third cost is short-termism. When a scheme measures its success by the number of AI models launched, it will keep backing projects that show results quickly and neglect research that takes years to pay off. The Noble Prize winning AlphaFold itself was trained on a database of protein structures that took half a century of largely public funding to build.

Doing it Better

Most of what is needed is not expensive. First, every faculty member at a recognised institution could get a small starter allocation, granted on a one-page request within a week and paid for in full, since a 40% subsidy is of little use to someone with no budget at all. Giving a thousand teachers 3,000 GPU hours each, at the full price of about Rs 110 an hour (the Rs 67 users pay plus the 40% subsidy), would cost about Rs 33 crore rupees a year. That is around 0.3% of the IndiaAI outlay, and less than what fifteen government bodies have already had paid for them in full.

Second, the NSM could set aside a share of its hours for outside users, especially at colleges and universities away from the major research institutions, with the same rules at every centre. Third, IndiaAI could name computational chemistry, physics, biology, materials and earth sciences in its researcher category, put scientists from these fields on its review panels, and publish who sits on them. Fourth, allocations could run for a year and be drawn on as needed, with unused time carried over. Fifth, NSM machines could serve as regional hubs for nearby universities, with support staff and ready-installed software, and the mission's training in high-

performance computing could be opened to college teachers. ANRF grants for computational projects could also come with NSM or IndiaAI time attached, so that winning a grant brings the computing it needs. Finally, both schemes could publish who gets their capacity, by type of institution, state and field, along with waiting times, and how many requests they turn down.

None of this needs a new mission. India rightly counts its petaflops and its GPUs, but it should also count how many of the researchers who could use them actually do. Today, a teacher with a PhD and a good research question, working far from the big institutes, still has no straightforward way to get the computing time needed to run a calculation. Fixing that would not cost much, and it would put a lot of training and talent back to work.

Sameer Saurav teaches chemistry at A.N.S. College, Nabinagar, Bihar. He studied at IISER Kolkata and has a PhD in computational chemistry from IIT Kanpur.