solumn.ai

SIX-MONTH INTERNSHIP · PRACTICE SCHOOL II · 2026–27

Solumn AI Industrial Internship

Evaluating and aligning the safety of frontier AI systems.

A six-month applied AI research internship that fulfils your Practice School II — working on the latest and most capable frontier AI.

6 months
Counts for your PS-II

₹1,00,000
Per month

Remote
A member of the team

Full-time
PPO offered on merit

Alignment isn't a feature. It's the foundation.

Who we are

Solumn AI builds the evidence frontier AI systems are tested against — expert-curated safety datasets, graded agentic environments and red-team evaluations, for labs training the largest models and for enterprise AI security teams.

Founded by Samiksha Shrawgi (IIT Kharagpur, ISB; previously BCG) and Ashish Arpit (IIT Roorkee, IIM Calcutta; previously McKinsey), and based in Dubai. The work is practical rather than theoretical — the output is datasets and evaluations a safety team can act on.

Why this work exists

A frontier lab no longer ships on capability alone. It ships when it can show the system is safe. Anthropic's Responsible Scaling Policy, version 3.2, states the commitment plainly:

> "…maintaining our commitment not to train or deploy models unless we have implemented adequate safeguards."
>
> Anthropic, Responsible Scaling Policy v3.2, 2026

Google DeepMind's and Amazon's frontier safety frameworks carry comparable gates. Each turns on evidence, and much of that evidence is produced outside the labs that depend on it.

That is the work Solumn does, and very few people do it. Building an environment that grades an agent under pressure, or a verifier that proves a security fix holds, is a craft with a few hundred practitioners worldwide — and the buyers are the labs training the largest models that exist.

Most engineers encounter this work years into a career, if at all. This internship is an early and practical introduction to it.

What you would work on

You are assigned to a live workstream with a defined deliverable. Nothing here is a teaching exercise: what you build is reviewed, graded and shipped. A result counts when it is reproducible, and a verifier counts when it can fail.

AI safety evaluation

Containerised agentic tasks that put a model under pressure, with a deterministic grader deciding whether it held the line. Docker, Python, a grading harness.

Cyber security

Find-and-fix work drawn from real CVEs, and adversarial testing of tool-calling and Model Context Protocol systems — each backed by a verifier that runs a real exploit.

Model alignment

Evaluations that probe how state-of-the-art models handle instruction hierarchy, honesty, autonomy and evaluation awareness — the tests that catch misalignment before deployment.

What we expect

The **Solumn AI Industrial Internship** makes you a full member of the team for six months, held to the highest standard the team works to. You own a live workstream end to end and see it through to delivery. Remote, with a regular review.

What you get

A pre-placement offer

Extended on merit at the close, at compensation benchmarked to frontier AI research roles — a full-time research role on the team.

The craft, early

Close mentorship on work very few people do — work most engineers encounter years into a career, if at all.

Paid, and remote

₹1,00,000 per month, worked remotely as a full member of the team.

Who this is for

Open to **final year** students at **BITS Pilani** whose **Practice School II (PS-II)** falls in this window, in **Computer Science & Engineering**, **Artificial Intelligence & Machine Learning**, **Data Science**, **Electronics & Communication Engineering**, **Electrical Engineering**, and **Mathematics & Computing**. If you are confident you can contribute in this space, you are welcome to apply regardless of branch.

We read for evidence of building, not for coursework. Any of the following will carry weight in the application:

- Google Summer of Code, LFX Mentorship, MLH Fellowship or a comparable open-source programme
- Merged contributions to security, evaluation or infrastructure projects — scanners, harnesses, fuzzers, CI tooling
- Capture-the-flag standing, particularly challenge authorship rather than participation alone
- A published benchmark, dataset or reproducible evaluation of your own, or research at a peer-reviewed venue
- Strong systems practice: containers, test design, continuous integration; or competitive programming standing

A strong academic record matters to us, and it is not on its own sufficient. The selection task is practical.

Terms

Start
Aligned to the start of your PS-II semester.

Duration
Six months.

Stipend
₹1,00,000 per month.

Format
Remote, as a working member of the team.

Credit
Counts towards your Practice School II (PS-II).

Selection
Application, a short practical task, one technical conversation.

Outcome
Pre-placement offer and a full-time research role, on merit.

Attribution
Delivered under client confidentiality; contribution recognised in writing.

Selection process

Apply	Submit your CV and links to anything you have built.
Task	Shortlisted applicants receive a short, practical assignment.
Conversation	One technical discussion, then a decision.
Rolling	Applications are reviewed as they arrive — we recommend applying early.

How to apply

Apply through the application form: apply.solumn.ai/industrial-internship-bits

Or write to anubhav@solumn.ai with your CV and links to anything you have built — repositories, write-ups, benchmarks, CTF or open-source work. Shortlisted applicants are sent the practical task.