

RACE WEEKLY INTELLIGENCE BRIEF


RETINA

AI · CYBERSECURITY · CLOUD · LEADERSHIP

THIS WEEK'S COVER:

AI STARTS INVENTING

For years, AI helped scientists work faster. Now it does the science itself. It designed a drug that is already being tested in people. It predicted the shape of almost every protein known to biology, and won a Nobel Prize doing it. In some labs it now runs its own experiments, day and night, without being asked. After a long run of covers about risk, this one is mostly good news. It is also a new kind of advantage, and most companies have no plan for it. Here is what changed, and what a leader does about it.



ALSO INSIDE

The World in 5 · Signal · Hyperfocal ·
Three Forces · Perspectives · The
Discovery Playbook · Metric of the
Week · Fun Fact · CANDID · Book
Shelf · Research Radar · RACE Pulse

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Discovery Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

FROM THE EDITOR'S DESK

For nineteen issues we asked what AI can do for people. This week the question changes. What happens when AI does the work itself.

In October 2024, two of the year's science Nobels went to AI. The Chemistry prize honoured the team behind AlphaFold, the system that cracked how proteins fold, a problem biology had chased for fifty years. The Physics prize went to the pioneers of the neural networks that make modern AI work. It was a quiet turning point. AI was no longer just a tool that scientists used. It had become an instrument that made discoveries of its own.

Since then the pace has been startling. A drug discovered and designed by AI, from a company called Insilico, has moved through early human trials and shown real improvement in lung function. Another AI-designed compound is now in late-stage trials. DeepSeek was the story of cheaper models. This is the story of what those models are starting to find. In some labs the whole cycle now runs on its own: the software proposes an experiment, a robot runs it, the software reads the result and proposes the next one. It does not tire, and it works through the night.

Two honest cautions keep this grounded. First, the lab has sped up, but the clinic has not. No AI-discovered drug has been approved yet, and most candidates still fail in humans for reasons no model can see in advance. Second, AI is not a genius in a box. It works inside problems that people define, and the real bottleneck now is testing ideas in the real world, not thinking them up. But the direction is clear, and it reaches far past pharma. Any business that explores a huge range of options slowly, by hand, is looking at its own version of this shift. India has placed its bet through the BioE3 policy and a network of Bio-AI hubs. This issue explains what changed, and the moves a leadership team can make now.

THE WORLD IN 5

GLOBAL HEADLINES A CXO SHOULD BE ABLE TO REFERENCE

- From the Editor's Desk
- **The World in 5**
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Discovery Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

Five signals from the past fortnight, worth a note at your next leadership table.

1. The year AI won the Nobel

In 2024, AI sat behind two science Nobels in the same week. The Chemistry prize went in part to the makers of AlphaFold for solving protein structure. The Physics prize went to the founders of neural networks. Prizes do not change the science, but they mark the moment a field stops arguing about whether something is real. AI as a scientific instrument is now the mainstream view, not a fringe claim.

2. A drug that AI designed is in human trials

Insilico Medicine's rentosertib, the first drug both discovered and designed by AI, cleared a Phase 2a trial in patients with a serious lung disease. Those on the drug gained lung function while the placebo group declined. A second AI-designed compound, from Schrodinger, is now in Phase 3. The caution matters as much as the news: no AI-discovered drug has been approved yet, and human trials remain the hard, slow part.

3. New materials by the hundred thousand

DeepMind's GNoME predicted 380,000 stable new materials in a single effort, more than researchers had catalogued in decades. Autonomous labs are starting to make and test them. Chemistry and materials lead here because the goal is often clear and measurable, which is exactly where the technology is strongest. One synthetic-biology firm has said it will move all of its research onto autonomous infrastructure by the end of 2026.

4. From tool to colleague

The newest systems do not just calculate. They propose. Google's AI co-scientist and Sakana's AI Scientist can generate hypotheses, design experiments, and write up results, running the research loop with light human oversight. The honest limit still holds. These systems shine where the question is well defined, and they still need people to judge messy results and to run the real-world test that settles the matter.

5. India: the policy and the talent, not yet the ownership

India has moved. The BioE3 policy backs biomanufacturing with a national network of Bio-AI hubs and biofoundries, and firms like Aganitha, working with CSIR-CCMB, are applying generative AI to therapeutic design. The risk is familiar. India has deep talent and runs a great deal of the world's AI research work, but the resulting intellectual property often sits abroad. The task is to own discoveries, not just service them.

Source:

<https://www.nobelprize.org/prizes/physics/2024/press-release/What/>
<https://insilico.com/news/xmjns4l091-insilico-initiates-phase-iii-clinical-trial>
<https://doi.org/10.1038/s41591-025-03743-2>
<https://www.takeda.com/newsroom/newsreleases/2026/zasocitinib-outperforms-deucravacitinib-study/>
<https://deepmind.google/blog/millions-of-new-materials-discovered-with-deep-learning/>
<https://research.google/blog/accelerating-scientific-breakthroughs-with-an-ai-co-scientist/>
<https://bmi.dbtindia.gov.in/>
<https://www.aganitha.ai/company/news-and-events/aganitha-and-ccmb-collaboration/>

 SIGNAL

THE INVENTION ENGINE, BY THE NUMBERS

- From the Editor's Desk
- The World in 5
- **Signal**
- Hyperfocal
- Panorama
- Perspectives
- The Discovery Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

2 Nobels

in 2024: AI shared the Chemistry prize (protein folding) and underpinned the Physics prize (neural networks)

~200M

protein structures AlphaFold predicted and released free, close to every protein known to science

380,000

stable new materials DeepMind's GNoME found in a single run, decades of discovery at once

~30 mo

from project start to first human trial for an AI-designed drug (Insilico), against the usual several years

+98 mL

lung-function gain at 12 weeks for that drug in its Phase 2a trial, while the placebo group declined

\$8.9B

size of the lab-automation market today, on track to \$24B by 2035 as self-driving labs scale

Source:

<https://www.nobelprize.org/prizes/physics/2024/press-release/What/>
<https://deepmind.google/science/alphafold/>
<https://deepmind.google/blog/millions-of-new-materials-discovered-with-deep-learning/>
<https://insilico.com/news/xmjnsn4l091-insilico-initiates-phase-iii-clinical-tr>
<https://doi.org/10.1038/s41591-025-03743-2>
https://ceder.berkeley.edu/publications/2023_Nathan_Alab.pdf



HYPERFOCAL

What Changed This Week

- From the Editor's Desk
- The World in 5
- Signal
- **Hyperfocal**
- Panorama
- Perspectives
- The Discovery Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

The Lab Became a Loop

Discovery used to run in one direction. A person had an idea. Months later, a slow experiment said yes or no. Then the person had another idea. The rate of progress was capped by how fast humans could think, and by how long each test took. That is the model AI is quietly dismantling.

The new shape is a loop. The software proposes a candidate, a molecule, a material, a design. A robot or a simulation tests it. The software reads the result and proposes a better one. The loop runs on repeat, at machine speed, often around the clock. This is what people mean by a self-driving lab. The key change is not that AI is smarter than a scientist. It is that the loop never stops, never forgets a result, and can explore a space far larger than any team could search by hand.

That moves the bottleneck. For a long time, ideas were scarce and testing was the easy part. Now capable AI can generate more good ideas than anyone can check, so the scarce resource becomes the experiment itself, and the data it produces. Two things follow for a leader, and they are worth saying plainly. The advantage goes to whoever can run the loop fastest, because speed of iteration is now speed of learning. And it goes to whoever owns the best private data, because public models are shared with everyone while your own results are not. The honest limit is real too. The loop works best when the goal is clear and measurable. It struggles with open-ended questions, it still needs people to read ambiguous results, and in the end the real world, a patient, a bridge, a market, is the judge that no model can replace.

RETINA TAKE

The story is not that AI replaced the scientist. It is that invention turned into a loop that runs without waiting for us. The winners will not be the ones with the cleverest single idea. They will be the ones who built the fastest loop on a problem they own, and fed it data no one else has.

PANORAMA

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- **Panorama**
- Perspectives
- The Discovery Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

THREE FORCES — AND THE MOVE EACH ONE DEMANDS

FORCE 1

Discovery Is Becoming a Search Problem

Finding a new drug, material or design means searching a space too large for people to explore. That is exactly the kind of problem AI is built for. AlphaFold searched the space of protein shapes. GNoME searched the space of stable materials. The lesson is general: when the answer is somewhere in a vast set of options, a machine can look in places a human team never could.

The Move: Take your hardest R&D or design question and ask whether it is really a search problem in disguise. If it is, it is a candidate for AI, and probably a strong one.

FORCE 2

The Experiment Is the New Bottleneck

Once AI can propose more good ideas than you can test, progress is limited by how fast you can run experiments and capture the results. This is why self-driving labs matter. They attack the slow, physical part of the loop, not the thinking part. The organisations pulling ahead are investing in fast, clean feedback, not just in bigger models.

The Move: Invest in the loop, not only the model. Automate testing and data capture so results flow back fast. Whoever iterates fastest learns fastest, and learning is the whole game.

FORCE 3


Your Moat Moves From Ideas to Data

The powerful models are increasingly shared. The frontier is available to your competitors too. What is not shared is your own experimental and operational data, the record of what you tried and what actually happened. In a world where everyone can rent the same intelligence, that private record becomes the real advantage, because it is the one thing a rival cannot buy or copy.

The Move: Treat your data as a strategic asset from today. Capture it cleanly, keep it, and protect it. It is the fuel your invention engine runs on, and no one else has yours.

Source:

<https://deepmind.google/science/alphafold/>
<https://deepmind.google/blog/millions-of-new-materials-discovered-with-deep-learning/>
https://ceder.berkeley.edu/publications/2023_Nathan_Alalab.pdf
<https://labs.cloudsecurityalliance.org/research/>



- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- **Perspectives**
- The Discovery Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs



PERSPECTIVES

DEEP TECH — ANALYSIS BY THE RACE TEAM

The Bottleneck Was Never the Ideas



Dr. Shinu Abhi
Professor and Director
Corporate Training, RACE,
REVA University

<https://www.linkedin.com/in/shinu-abhi-phd-711a339>

The most useful way to read AlphaFold is not as a miracle. It is proof that discovery can be turned into an engineering discipline. You define the space to search. You build a loop that proposes and tests. You feed it good data. That recipe is not unique to protein folding. It is starting to work across chemistry, materials, and drug design, and it will keep spreading to any field where the question can be made precise.

This reframes what a research advantage even is. For most of history, the scarce thing was the brilliant idea. That is no longer the constraint. Capable AI can generate more

plausible ideas than any lab can test. So the advantage shifts to two things that are unglamorous but decisive: how fast you can run the test-and-learn loop, and how much of your own high-quality data you have kept. Neither is about having the smartest people in the room. Both are about how you have chosen to build and instrument your work.

For Indian organisations, I see a clear opening and a clear risk, side by side. We have the talent, and through efforts like BioE3 we now have the policy. But we have long been comfortable running other people's research while the discoveries, and the value, settle somewhere else. The discovery loop does not require a billion-dollar lab. It requires a real problem you own, a tight loop, and data you have gathered with care. That is well within reach here. We teach it as frugal AI, and it is the difference between servicing the future and owning a piece of it.

RACE Perspective

AI has turned invention into a loop anyone disciplined can run. You do not need to build the frontier. You need one problem that is truly yours, a fast cycle of propose, test and learn, and data no one else holds. Own the problem and the data, rent the rest.



- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- **The Discovery Playbook**
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs



THE DISCOVERY PLAYBOOK

PUTTING THE
INVENTION ENGINE TO WORK

Five moves, ordered from fastest to most strategic, that any leadership team can act on.

MOVE 1

FIND YOUR SEARCH PROBLEM

This is not only for pharma. Almost every business explores a large space slowly and by hand. Recipes and formulations. Materials and suppliers. Product designs. Routes and schedules. Pricing and promotions. Somewhere you are testing options one at a time when a machine could search thousands. Name that space.

The Move: Pick one high-value area where you currently explore options by hand, and ask a simple question. Could AI search this space far wider and faster than we do now.

MOVE 2

BUILD THE LOOP, NOT JUST THE MODEL

The advantage is iteration speed, not the model you pick. A good loop proposes, tests, reads the result, and proposes again, with as little friction as possible. Most organisations lose weeks moving results by hand between steps. Close that gap and the whole cycle speeds up.

The Move: Instrument the feedback loop so results flow back automatically. Measure how long one full cycle takes today, then work to cut it. Cycle time is your real metric.

MOVE 3

GUARD YOUR DATA LIKE THE ASSET IT IS

When everyone can rent the same models, your own data is the edge. The record of what you tested and what happened is something no competitor can copy. Yet in most companies it is scattered, unlabelled, or thrown away. That is a strategic asset being treated as exhaust.

The Move: Start capturing your experimental and operational data cleanly now, and keep it. You will want years of it later, and you cannot go back and collect the past.

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- **The Discovery Playbook**
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs



THE DISCOVERY PLAYBOOK

SPENDING ON AI WITHOUT
BETTING ON THE BUBBLE

MOVE 4

KEEP PEOPLE ON JUDGMENT, MACHINES ON VOLUME

The loop is superb at more and faster. It is weak at deciding what matters and reading a strange result. Those are human jobs, and they get more important, not less, as the volume rises. The failure mode is automating the routine and then letting the machine make calls it is not ready to make.

The Move: Automate the repetitive middle of the loop, and reserve your best people for setting the goal and judging the hard results. Do not blur that line.

MOVE 5

RENT THE FRONTIER, OWN THE PROBLEM

You do not need your own AlphaFold. The frontier models, cloud labs and specialist tools already exist and improve every month. Building your own from scratch is usually the expensive, slow path. What you should never outsource is the problem itself and the data it generates, because that is where your advantage lives.

The Move: Buy or partner for the general capability, and keep the problem, the data and the judgment in house. Build only the part that is genuinely yours.

Source:

<https://deepmind.google/blog/millions-of-new-materials-discovered-with-deep-learning/>
<https://insilico.com/news/xmjns4l091-insilico-initiates-phase-iii-clinical-tr>
<https://research.google/blog/accelerating-scientific-breakthroughs-with-an-ai-co-scientist/>
<https://race.reva.edu.in/>

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Discovery Playbook
- **Metric Of The Week**
- **Fun Fact**
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

METRIC OF THE WEEK

ONE NUMBER, PROPERLY UNDERSTOOD

Zero **AI-discovered drugs approved by regulators so far. The lab got fast. The clinic did not.**

What it is: Despite billions invested and a growing pipeline, not a single drug discovered by AI has yet been approved as of 2026. AI has compressed the early work dramatically, some candidates reached human trials in half the usual time, but roughly nine of ten drug candidates still fail once they reach people, for reasons of safety or effectiveness that show up only in real bodies.

Why it matters now: It separates the part of discovery AI has changed from the part it has not. Searching for a candidate is faster. Proving it works in the real world is not. Confusing the two is how a leader overpays for a breakthrough that is still years from paying off.

How leaders use it: Put AI to work where the feedback loop is fast and cheap, which is early search and design. Stay patient and

disciplined where reality is the judge, which is late-stage proof. Fund the first aggressively. Do not pretend the second has changed.

Source:

<https://doi.org/10.1038/s41591-025-03743-2>

<https://www.nature.com/articles/s41573-026-01496-2>

<https://intuitionlabs.ai/articles/ai-discovered-drugs-clinical-trials-2026>

FUN FACT

In 2025, a research team fed a hard question to Google's AI co-scientist: how do certain bacteria spread their resistance to drugs. In about two days, the system proposed the same core answer the scientists had spent years working out in the lab, an answer they had not yet published, so the AI could not have read it anywhere. The researchers were, by their own account, stunned. The machine did not do the experiments, and it did not replace the decade of human insight behind the question. But it compressed the hardest part of thinking, the leap to the right hypothesis, from years into an afternoon.

Source: Google Research AI co-scientist (2025); reporting on the Imperial College / bacterial gene-transfer study

SOMETHING TO REPEAT AT DINNER

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Discovery Playbook
- Metric Of The Week
- Fun Fact
- **Candid**
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs



CANDID



Usha Rengaraju

AI Engineer, Kaggle Grandmaster

<https://www.linkedin.com/in/usha-rengaraju-b570b7a2>

Q1. AI now designs drugs that reach human trials and predicts materials by the hundred thousand, yet no AI-discovered drug has been approved. How should a board separate real acceleration from headline acceleration?

Notice that AI accelerates the cheap front of the pipeline (generating candidates) while value sits at the back (surviving trials), and no AI-discovered drug is FDA-approved yet, with the leading case only reaching Phase III in mid-2026. A board should ignore output counts and ask whether AI relieved the binding constraint; the real metrics are survival rate at each gate versus a matched baseline, cost per validated result, and a rising early-kill rate (failing faster and cheaper is the win). Headline acceleration counts proposals; real acceleration shows more things surviving the gate that actually kills you.

Q2. The bottleneck in discovery is shifting from ideas to experiments and to proprietary data. For a company that is not in pharma or materials, is there still a version of this shift worth acting on, and where would you start?

Yes, the shift is universal: when ideas become cheap, advantage moves to experiment throughput and proprietary outcome data. Your "lab" is any

decision you make thousands of times where results are fast and reversible, pricing, targeting, routing, retention. Start by instrumenting one such decision so every action produces labelled outcome data, then use AI to generate candidate policies, test them, and feed results back. The durable asset isn't the model; it's the experiment platform plus proprietary feedback no competitor or foundation model has seen.

Q3. Self-driving labs run the propose, test and learn loop around the clock. What has to be true inside an organisation for that loop to actually pay off, rather than just automate busywork faster?

The loop pays off only if it sits on the true bottleneck (not a cheap step whose queue just grows), genuinely closes so "learn" changes the next batch with no human handoff, optimises a real objective the machine can read honestly rather than a convenient proxy, and runs where experiments are cheap, and feedback generalises. Organisationally, hold people accountable for validated discoveries, not experiments-per-week, capture results in machine-readable form, and redeploy freed staff to hard judgment work, otherwise you've just built a very fast busywork generator.

Q4. India has the policy and the talent, but often ends up running other firms' AI while the intellectual property sits abroad. What would it take for Indian enterprises to own discoveries rather than service them?

The gap is capital and incentives, not talent: owning discovery means holding a scarce upstream asset, proprietary data, foundational IP, or infrastructure, whereas services supplies the labour, the one part that never accrues equity. India's edge is data it generates by right (population-scale health, agriculture, twenty-plus languages, DPI exhaust),



CANDID

CXO VOICE — IN CONVERSATION THIS WEEK

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Discovery Playbook
- Metric Of The Week
- Fun Fact
- **Candid**
- Book Shelf
- Research Radar
- RACE Pulse
- RACE Programs

which must become home-built platforms rather than be shipped abroad. It would take patient risk-tolerant capital, anchor procurement that favours owning over renting, retention of senior researchers who can lead discovery, and an honest choice of vertical battlefields, health, agri, Indian-language AI, over competing head-on with frontier labs. Services is a comfortable local optimum; the shift is an appetite change before a capability one.

Note: CANDID features direct contributions from CXO leaders.
To contribute, contact: retina@reva.edu.in



BOOK SHELF

RETINA'S READ OF THE WEEK

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Discovery Playbook
- Metric Of The Week
- Fun Fact
- Candid
- **Book Shelf**
- Research Radar
- RACE Pulse
- RACE Programs

RETINA RECOMMENDS · THIS WEEK'S THEME

The Coming Wave: Technology, Power, and the Twenty-first Century's Greatest Dilemma


Mustafa Suleyman · Crown, 2023

Suleyman co-founded DeepMind, the lab behind AlphaFold, so he writes about AI in science from the inside. His central claim fits this issue exactly. The real wave is not AI alone. It is AI joined to biology and to the tools of invention, arriving together and faster than our institutions expect. Read alongside this week's cover, the book explains why a drug designed by software and a Nobel Prize for protein folding are not separate news items. They are the same wave, breaking early.

What it gives a leadership team is a sense of scale and of pace. Suleyman is not a cheerleader. He spends much of the book on the harder half of the story, how technologies this powerful spread, and

how hard they are to contain once they do. For a CXO, that balance is the value. It makes the opportunity feel real without making it feel safe, which is the correct posture for a technology that both invents medicines and lowers the cost of causing harm.

It is not a manual, and it will not hand you a strategy. What it offers is judgment: a way to think about a wave that will touch your industry whether or not you work in science. Read it before you decide how much of your R&D, and your caution, this moment deserves.



- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Discovery Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- **Research Radar**
- RACE Pulse
- RACE Programs



RESEARCH RADAR

RACE PUBLICATIONS THIS WEEK

EDITOR'S NOTE

Two new publications from RACE. One keeps cloud systems standing when a breach hits. One keeps trains on the tracks. Standard format: a concise technology summary followed by clear enterprise relevance.

Cloud Security · Containers · Incident Response and Recovery

Cyber Resilience for Containerized Workloads: A NIST-Based Approach to Incident Management and Recovery

Saritha Nagaraju; Nishanth Kumar Pathi; Rashmi Agarwal

Containers now run much of the enterprise stack, but their short life makes them hard to secure. A container can spin up, do its work, and vanish before anyone can watch it, detect a threat, or preserve evidence, and standard incident-response tools were not built for that pace. This work proposes an end-to-end framework grounded in the NIST Cybersecurity Framework, pairing runtime detection (Falco) and host monitoring (Wazuh) with automated recovery (ArgoCD and Terraform), so a compromised service can be caught, investigated, and restored with minimal downtime. In testing it improved both detection and recovery times and eased forensic evidence collection.

Enterprise Relevance: For any organisation running Kubernetes or Docker in production, this answers a real gap: how to keep incident response and forensics working when the thing you are defending disappears in seconds. Anchoring the design in the NIST framework gives security and compliance teams a shared, auditable vocabulary, and recovery-by-automation turns a breach from an outage into a fast, routine restore.

Proceedings of Fifth Emerging Trends and Technologies on Intelligent Systems (ETTIS 2025) · Lecture Notes in Networks and Systems, vol 1591, pp 24–34 · Springer, Singapore, 2026

https://link.springer.com/chapter/10.1007/978-981-95-0681-1_3

Computer Vision · Deep Learning · Rail Infrastructure Safety

A Deep Learning System for Comprehensive Railway Track Fitness Assessment and Defect Identification

M. Bharath; Shinu Abhi

Inspecting rail track by hand is slow, costly, and dangerous, and every hour a fault goes unseen is an hour closer to a derailment. This work puts the inspection on the train itself. Using video and images captured as the train runs, YOLOv8 models flag critical defects in real time: broken rails, cracked sleepers, damaged fishplates, faulty fasteners, and gauge variations. The models were trained on annotated data from multiple locations and lighting conditions so detection holds up in the real world, not just the lab, cutting cost and human error while catching problems sooner.

Enterprise Relevance: Beyond railways, this is a template for turning any linear physical asset, track, pipeline, power line, roadway, into something a camera and a model can inspect continuously instead of by periodic manual survey. The pattern is the same: move inspection from scheduled and human to constant and automated, catch failures earlier, and lower both the cost and the safety risk of finding out too late.

2025 Eighth International Conference on Image Information Processing (ICIIP), Solan, Himachal Pradesh · IEEE (added to IEEE Xplore, January 2026)

<https://ieeexplore.ieee.org/document/11346317>

Explore the full RACE research archive at <https://raceresearch.reva.edu.in>. For collaboration, licensing or research-partnership enquiries: race@reva.edu.in

RACE PULSE

EVENTS · COMMUNITY · WHAT'S COMING

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Discovery Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- **RACE Pulse**
- RACE Programs

RACEx360 SEASON 2

25 SEPTEMBER 2026 · TAJ WEST END, BENGALURU

One day, one stage, four frontiers.

REVA Academy's Emerging Technology Conference returns, in partnership with The Economic Times — 750+ builders, buyers and decision-makers across four tracks that matter: GenAI & Agentic AI · Cloud & Platform Engineering · Cybersecurity in the AI Era · Deep Tech & Sustainability.

Not a day of slides — a day of proof: an Innovation Experience Zone with live demos across AI, robotics, cybersecurity, cloud and quantum; live build sessions; an AI Hackathon and Cybersecurity Challenge; and keynotes that close with a 90-day action plan you can take back to your team.

The evening turns gold with the Top 10 Women Tech Leaders India Awards 2026, now in its 4th edition — India's only annual awards programme dedicated to women technology leaders.

Early Bird 4,000 + GST until 15 August · Standard 5,000 + GST. Passes cover all four tracks, the Expo Hall and the Experience Zone.

GET YOUR PASS:

<https://race.reva.edu.in/racex360-season-2>

BECOME A SPONSOR:

<https://race.reva.edu.in/racex360-sponsorship>



RACEx360
Emerging Technology Conference 2026



In partnership with **ET THE ECONOMIC TIMES**

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Discovery Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- **RACE Pulse**
- RACE Programs

RACE PULSE

EVENTS · COMMUNITY · WHAT'S COMING

NAAC
GRADE **A+**
ACCREDITED UNIVERSITY

REVA
RACE
REVA Academy for
Corporate Excellence

THE ENTERPRISE

GAVE AI THE KEYS

You get 90 minutes to prove it shouldn't have.

EC-Council

HACKai

Challenge

Prompt injection. RAG poisoning. Jailbreaking.
Agentic privilege escalation.

QUALIFIER

10th to 18th September 2026

FINALE

25th September 2026 at
The TAJ West End, Bengaluru

CASH PRIZE
₹1.75 Lakhs
+
CEH Vouchers

REGISTER TODAY

Registration closes on 10th September 2026

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Discovery Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- **RACE Pulse**
- RACE Programs

RACE PULSE

EVENTS · COMMUNITY · WHAT'S COMING

OCTOBER 2026 BATCH - APPLY BY 10 SEPTEMBER 2026

Applications open for all six RACE Master's & PG Programmes — Business Analytics (#1 India), Artificial Intelligence (#3 India), Cybersecurity (#5 India), and Cloud Architecture & Security.

Weekend format. AICTE/UGC approved. Industry mentors from Amazon, Microsoft, AWS, EC-Council, Terralogic, IBM and more.

To know more, call:
+91 89040 58866

enquiry@race.reva.edu.in ·
Apply now: race.reva.edu.in

NEW BATCH **Certified Agentic AI Engineer** 26 SEPTEMBER 2026

Practical expertise in LLMs, AI agents, multi-agent systems, agentic workflows, governance and capstone-based implementation. Real-world use cases, deployable solutions.

To know more, call:
+91 89040 58866

Apply now:
<https://race.reva.edu.in/certified-agenticai-engineer>

RACE SANDBOX SATURDAY 26 SEPTEMBER 2026

Run your career simulation, live on campus. Sit in a live Master's class, walk the AI, Cybersecurity and Cloud labs, watch participants and alumni demonstrate real projects, and take a 1:1 on where your career goes next. Experience today, decide your tomorrow.

9:30 AM – 2:00 PM
RACE, REVA University,
Bengaluru

Reserve your slot:
+91 89040 58866 ·

enquiry@race.reva.edu.in ·
<https://race.reva.edu.in/>

- From the Editor's Desk
- The World in 5
- Signal
- Hyperfocal
- Panorama
- Perspectives
- The Discovery Playbook
- Metric Of The Week
- Fun Fact
- Candid
- Book Shelf
- Research Radar
- RACE Pulse
- **RACE Programs**

Master's & PG Programmes

New Batch: September 2026 · 12–24 months ·
Weekend · For Working Professionals

#1 IN INDIA · ANALYTICS INDIA MAGAZINE

ANALYTICS

PG Diploma / M.Sc. in
BUSINESS ANALYTICS
AWS Academy · Microsoft Azure

12–24 Months

#3 IN INDIA · TECHGIG

ARTIFICIAL INTELLIGENCE

M.Tech. in
ARTIFICIAL INTELLIGENCE
AWS Academy · Microsoft Azure

24 Months

#3 IN INDIA · TECHGIG

ARTIFICIAL INTELLIGENCE

PG Diploma / M.Sc. in
ARTIFICIAL INTELLIGENCE
AWS Academy · Microsoft Azure

12–24 Months

#5 IN INDIA · ANALYTICS INDIA MAGAZINE

CYBERSECURITY

M.Tech. in
CYBERSECURITY
EC-Council · Microsoft Azure · Terralogic

24 Months

#5 IN INDIA · ANALYTICS INDIA MAGAZINE

CYBERSECURITY

PG Diploma / M.Sc. in
CYBERSECURITY
EC-Council · Microsoft Azure · Terralogic.

12–24 Months

CLOUD

PG Diploma / M.Sc. in
**CLOUD ARCHITECTURE &
SECURITY**
AWS Academy · Microsoft Azure

12–24 Months

Short-Term Certification Programmes

100 Hours · Project-Driven · Customisable for Organisations

AGENTIC AI | 4–6 weeks

AGENTIC AI SYSTEMS DESIGN

Multi-agent systems using LangChain, LangGraph, CrewAI. Capstone:
deployed autonomous agent.

AI ENGINEER | 8–12 WEEKS

AI ENGINEERING CERTIFICATION

Data pipelines, model training, LLMs, deployment. Real production
projects every batch.

CLOUD | 8–12 WEEKS

CLOUD ARCHITECTURE & DEVSECOPS

Multi-cloud, IaC, secure CI/CD on real AWS/Azure/GCP environments. Not
demos.

CYBERSECURITY | 8–12 WEEKS

CYBERSECURITY & SOC OPERATIONS

Threat detection, incident response, SIEM/SOAR. Hands-on in RACE's live
Security Operations Centre.

DATA & AI | 12 WEEKS

DATA SCIENCE & BUSINESS ANALYTICS

Statistics to ML to BI. Python, SQL, Power BI, Tableau on real business
datasets.

LEADERSHIP | 8–12 WEEKS

TECHNO-FUNCTIONAL LEADERSHIP

Executive presence, decision-making, leading tech-driven teams. Built for
mid-to-senior managers.

Enroll with us · race@reva.edu · race.reva.edu.in

RETINA

AI · CYBERSECURITY · CLOUD · LEADERSHIP

THIS WEEK, ASK YOUR LEADERSHIP TEAM

“Of every rupee we’ll spend on AI this year, how much is tied to a business number we’re already tracking, and if the buildout corrects, which of our tools and vendors would we lose?”

REVA Academy for Corporate Excellence,
REVA University
Rukmini Educational Charitable Trust
Kattigenahalli, Yelahanka, Bangalore, Karnataka, India, 560 064
Phone: +91 89040 58866

RETINA PUBLISHES EVERY SUNDAY · FORWARD TO A CXO YOU RESPECT

