UPSC & STATE PCS CURRENT AFFAIRS · UJIYARI.COM

EDITORIAL ANALYSIS

The Illusion of Wisdom: Why a Process Can Be Automated but Judgement Cannot

MINT

27 July 2026 · SCIENCE & TECH · ECONOMY · GS3 · GS4

CURATED & WRITTEN BY

Bharat Choudhary

UPSC Educator & Content Creator

linkedin.com/in/epicbharat

The Illusion of Wisdom: Why a Process Can Be Automated but Judgement Cannot

Mint 27 July 2026 GS3 GS4

INTERVIEW ANGLE

"When an algorithm advises a public official and the decision goes wrong, who is morally responsible: the official who followed it, the agency that procured it, or nobody at all?"

Every fact web-verified against primary sources

THE LIFT LINE

A process can be automated. Wisdom cannot. Ford Motor Company spent years widening the use of artificial intelligence in engineering quality control, then brought hundreds of veteran engineers back because the automated inspection systems were not delivering, and the retreat is more instructive than the advance. The employment question, whether AI takes jobs, is the shallow one; the deeper and less appreciated risk is epistemic, because AI is changing not only how we work but how we decide what to believe.

WHY THIS EDITORIAL MATTERS FOR YOUR EXAM

GS Paper 3 owns the technology and economy half. Named syllabus hooks: “awareness in the fields of IT, space, computers, robotics”; “developments and their applications and effects in everyday life”; and “achievements of Indians in science and technology, indigenisation of technology”. **GS Paper 4** owns the harder half: moral agency and accountability when a machine advises, “ethical concerns and dilemmas in government and private institutions”, “laws, rules, regulations and conscience as sources of ethical guidance”, and probity in governance where automated systems mediate entitlements.

For **Prelims**, hold the institutional map: the IndiaAI Mission and its approved corpus, the India AI Governance Guidelines of 5 November 2025 with the AI Governance Group, the Technology and Policy Expert Committee and the AI Safety Institute, the Digital Personal Data Protection Act, 2023 and its Rules of 13 November 2025, NITI Aayog’s Responsible AI for All approach documents, and the fact that India has no dedicated AI statute. For **Mains**, the transferable distinction is between information, knowledge and judgement, and the proposition that automation transfers the easy part of a task to a machine and leaves the hard part to a human who has lost the practice.

BACKGROUND AND CONTEXT

The trigger is corporate. In late June 2026 Ford Motor Company confirmed it was bringing back hundreds of experienced technical specialists, reported at around 350 and referred to internally as “gray beard” engineers, to strengthen product quality after years of expanding automated engineering and inspection systems that failed to deliver the expected results. Chief Operating Officer **Kumar Galhotra** said the company “had been relying more and more on automated quality systems and not getting the desired results”, adding: “We brought back technical specialists, and they hunt for failure points before a part ever reaches the plant floor.” The diagnosis reported alongside is the important part: experienced engineers had left before their tacit knowledge could be encoded into Ford’s systems, so the automated tools amplified weak inputs rather than catching design flaws.

The Mint column of 27 July 2026 by Saumitra Bhaduri reads this as an instance of what he calls the **illusion of wisdom** that AI creates, and reaches for a precedent. In the 1990s organisations rushed to embrace enterprise software and workflow automation, believing that digitising existing processes would automatically improve productivity. Many simply automated their own inefficiency. The management thinker **Michael Hammer** condensed the lesson into the title of his *Harvard Business Review* article of July to August 1990, “Reengineering Work: Don’t Automate, Obliterate”: technology should redesign work rather than merely accelerate flawed processes. Bhaduri argues that AI risks repeating that mistake on a far greater scale, and that public debate has focused on whether AI will eliminate jobs while missing the deeper and less appreciated risk, which is epistemic.

WAVE	PERIOD	THE PROMISE	THE FAILURE MODE
Enterprise software and workflow automation	1990s	Digitising a process raises productivity	Automating inefficiency; Hammer’s “Don’t Automate, Obliterate”
Statistical process control and robotics	2000s to 2010s	Machines catch what humans miss	Monitoring burden shifts to humans; vigilance decays
Generative and predictive AI	2020s	Systems that reason, not just execute	Fluent output mistaken for verified knowledge

THE CORE ARGUMENT / ISSUE

The author’s thesis: automation is not wisdom

Bhaduri’s claim is narrow and defensible. A process is a sequence of steps with a known structure, and structure is exactly what machines automate well. Wisdom is the capacity to recognise when the structure does not apply, which is the residual the automation leaves behind. Ford’s admission is a clean example: the

inspection process was automated, the judgement that decides which failure modes are worth hunting for was not, and removing the people who held that judgement removed the very input the automated system depended on.

The epistemic risk, as analysis rather than as the author's argument

The supplied text of the column ends at the framing of the epistemic risk, so what follows is Ujiyari's own reading of the research literature rather than an extension of Bhaduri's case.

Automation bias and automation complacency are the two documented human factors failures. Automation bias is the tendency to accept a system's recommendation over contradictory evidence; complacency is the reduced monitoring that follows trust in a reliable system. Both are strongest exactly where the system is usually right, because a system that fails often is watched, and a system that fails rarely is not.

The irony of automation, named by Lisanne Bainbridge in "Ironies of Automation" in *Automatica* in **1983**, is the structural version. Automating the routine parts of a task leaves the operator responsible for the rare, difficult residue, while removing the daily practice that would have kept those skills sharp. Bainbridge's conclusion inverts the intuitive one: automation increases rather than reduces the training an operator needs, and it converts active work into exhausting vigilance. The paper remains among the most cited in the history of human factors engineering.

Aviation supplies two verified case studies that point in opposite directions. **Air France Flight 447** crashed over the Atlantic on 1 June 2009; the final report of the French Bureau of Enquiry and Analysis for Civil Aviation Safety, published on 5 July 2012, found that iced pitot probes produced inconsistent airspeed data, the autopilot disconnected as designed, and the crew, whose training had not prepared them for manual handling at high altitude, never understood that they were in a stall and so never applied a recovery manoeuvre. The **Boeing 737 MAX** failed the opposite way: the Manoeuvring Characteristics Augmentation System (MCAS) acted on a single angle-of-attack sensor without cross-checking the second, and erroneous readings triggered it in Lion Air Flight 610 on 29 October 2018 and Ethiopian Airlines Flight 302 on 10 March 2019, killing 346 people in total and grounding the fleet worldwide. AF447 shows what happens when humans lose the skill the machine used to exercise; the MAX shows what happens when a machine is given authority without redundancy. Both are automation-dependency failures, and any serious framework has to guard against both directions at once.

Model collapse is the newest failure mode and the most directly epistemic. Shumailov and colleagues showed in *Nature* in **2024** that generative models trained recursively on model-generated data degrade in a characteristic way: distributional errors accumulate first and the model drifts from the true distribution, and then low-frequency events, the tails of the original content distribution, disappear permanently. As synthetic text saturates the open web, the practical implication is that fresh human-generated data must be periodically injected, and that the rare case, which is precisely the case judgement exists to handle, is the first casualty.

The **homogenisation** problem sits on top of this. When a large share of professionals consult the same handful of models, the diversity of options actually considered narrows even where no individual output is wrong. An epistemic monoculture is fragile in the same way an agricultural monoculture is: a single systematic

error propagates everywhere at once, and there is no independent stock of judgement left to catch it.

LAYER	WHAT THE MACHINE SUPPLIES	WHAT IT CANNOT SUPPLY
Information	Retrieval, aggregation, fluent summary	Whether the source is trustworthy
Knowledge	Pattern extraction, correlation, prediction	Whether the pattern holds outside the training distribution
Judgement	Ranked options, confidence scores	Which value to prefer when options conflict, and who answers for it

The Indian institutional position

India has chosen a light-touch, principles-first route. The **IndiaAI Mission** was approved by the Union Cabinet in March 2024 with a five-year corpus of **Rs 10,371.92 crore**. The **India AI Governance Guidelines** were released by the Ministry of Electronics and Information Technology (MeitY) on **5 November 2025**, creating an AI Governance Group, a Technology and Policy Expert Committee and an AI Safety Institute, and relying on existing law and voluntary safeguards rather than a new statute. NITI Aayog's **Responsible AI for All** approach documents, beginning in February 2021, set out principles of safety and reliability, equality, inclusivity and non-discrimination, privacy and security, transparency, accountability and the protection of positive human values, and recommended that sectoral regulators translate them into sector-specific rules. The **Digital Personal Data Protection Act, 2023** is the nearest binding instrument; its Rules were notified on **13 November 2025** with a phased eighteen-month compliance runway, consent manager registration due from November 2026 and most operational provisions from 2027. There is still no dedicated AI statute.

The Supreme Court of India offers the useful counter-model on deployment. SUVAS, the Supreme Court Vidhik Anuvaad Software launched in November 2019, translates judicial documents between English and Indian languages, and SUPACE, the Supreme Court Portal for Assistance in Court's Efficiency launched on 6 April 2021, assists with the factual and documentary complexity of case files. Both were framed explicitly as assistive, with the then Chief Justice of India stating that AI would not do decision making and that judges' autonomy and discretion would be retained. The boundary was drawn at the point where the machine would begin to exercise judgement rather than support it.

HOW TO THINK ABOUT THIS (ANALYTICAL FRAME)

Carry in **the Four Questions Before You Automate**.

1. **Obliterate or accelerate?** Hammer's question. Is this process worth having at all, or are we about to make a bad process faster? Automating an unexamined workflow locks it in and makes it harder to reform later.

2. **Where does the residual judgement sit?** Bainbridge's question. Name the human who handles the case the system cannot, and ask honestly whether that person still practises the skill often enough to exercise it under pressure.
3. **What breaks when the input degrades?** The Ford and model collapse question. Trace the system's dependence on fresh, high-quality human input, and ask what happens when the people who generated that input are gone.
4. **Who answers for the outcome?** The GS4 question. A machine cannot be a moral agent, because it cannot be held to account, cannot be asked for the reasons it actually used, and cannot bear consequences. Accountability must therefore attach to a named human or institution before deployment, not be improvised after failure.

The counter-argument, stated fairly. Automation delivers real, measurable gains, and romanticising human judgement is its own error. Human decision-makers are themselves biased, inconsistent, fatigue-sensitive and slower; automated systems have raised throughput and cut error rates across translation, logistics, fraud detection and diagnostic triage. Ford itself did not abandon AI; the reported strategy is to pair veteran engineers with the automated systems and use their expertise to retrain the models on better inputs, which is augmentation rather than reversal. Nostalgia for the pre-automation workplace also tends to forget the tedium and the error rate that the automation displaced. The honest position is not that AI degrades judgement, but that judgement is a distinct capability with its own maintenance cost, and organisations that stop paying that cost discover the bill later, at Ford's price or at AF447's.

THE DIAGRAM IN WORDS

Picture a task as a pyramid. The wide base is routine execution, the middle band is pattern recognition, and the narrow apex is judgement under ambiguity. Automation eats upward from the base, absorbing execution first and then much of pattern recognition, and it does so efficiently. But the apex does not shrink; it simply becomes a larger share of what remains for the human, who now reaches it less often and with less practice, because the base that used to train them has been removed. Now draw two arrows out of the apex. One points back down into the system as training data and design input, which is the arrow Ford discovered it had cut. The other points outward to a named accountable person, which is the arrow governance has to draw explicitly, because no algorithm draws it on its own. Cut the first arrow and the system degrades; leave the second undrawn and nobody answers when it does.

WAY FORWARD

1. **Mandate a named accountable officer for every automated decision system in government.** Adopt the human-in-command principle from the India AI Governance Guidelines as a binding procurement condition: no public system that determines an entitlement, a benefit or a penalty should go live without an identified official empowered to override it and answerable for its outputs.

2. **Build skill retention into automation contracts.** Following Bainbridge, require periodic manual-mode practice, scenario drills and shadow review for operators of high-consequence automated systems, and budget for it, so that the residual judgement the design depends on is still available when the exception arrives.
3. **Protect fresh human data as critical infrastructure.** Guard against model collapse by requiring provenance labelling of synthetic content, which the Information Technology Amendment Rules notified on 20 February 2026 have begun for synthetically generated information, and by funding curated, consented Indian-language human corpora under the IndiaAI Mission.
4. **Legislate an exclusion audit for welfare automation.** Automated eligibility, biometric authentication and de-duplication in schemes such as the Public Distribution System and rural employment guarantee attendance systems should carry a statutory duty to publish false-negative and exclusion rates disaggregated by district, with a guaranteed non-digital fallback, because in welfare delivery a false negative is not an error rate but a household without food.
5. **Move from principles to enforceable duties in stages.** Retain the light-touch posture for low-risk applications, but convert the NITI Aayog principles into binding obligations for high-risk public-sector uses, using the Digital Personal Data Protection Act, 2023 and its 2025 Rules as the compliance backbone rather than waiting for a standalone AI statute.
6. **Follow the judiciary's assistive model across public administration.** SUVAS and SUPACE were adopted with an explicit limit at the point where the machine would begin to decide rather than assist. Executive agencies deploying AI should state the same limit in writing, in the tender document, not in a press release after deployment.
7. **Institutionalise contestation to prevent epistemic monoculture.** Require public agencies procuring AI systems to document the alternatives considered and to preserve at least one independent review path outside the model, so that a single systematic model error cannot propagate unchallenged across an entire administrative function.

PYQ LINKAGE AND PRACTICE

UPSC has approached this theme from both papers. GS Paper 3 (2020) asked: "What are the main socio-economic implications arising out of the development of IT industries in major cities of India?" and GS Paper 3 (2018) asked candidates to discuss the potential of the fourth industrial revolution and its impact on employment in India. GS Paper 4 has repeatedly tested accountability and moral responsibility, including the 2019 question on the ethical issues involved in the use of social media, and case studies in which an official must decide whether to follow a rule or to exercise conscience. The bridge this editorial builds, that a machine can hold a rule but not a conscience, is exactly the ground where GS3 and GS4 meet.

Practice question: "Artificial intelligence can automate a process but not the judgement that gives the process meaning." Examine this proposition with reference to the deployment of automated decision systems in Indian public administration, and suggest an accountability framework. (250 words, 15 marks)

Interview angle: When an algorithm advises a public official and the decision goes wrong, who is morally responsible: the official who followed it, the agency that procured it, or nobody at all?

Sources: Mint, PIB, Ministry of Electronics and Information Technology, NITI Aayog, IndiaAI

Source: The Illusion of Wisdom: Why a Process Can Be Automated but Judgement Cannot — Ujiyari.com | Free UPSC & State PCS Editorial Analysis

RELATED DAILY ARTICLES

27 Jul Current Affairs Today, July 27, 2026

27 Jul BHAVYA-Rasayan Scheme: India Bets Rs 3,030 Crore on...

26 Jul The First Gatekeeper Penalty: EU Fines Google 890...

26 Jul No Gold Standard: The US-Saudi Civil Nuclear Deal and...

CURATED & WRITTEN BY

Bharat Choudhary

UPSC Educator & Content Creator

linkedin.com/in/epicbharat

📖 Read Full Article on Ujiyari →

https://ujiyari.com/editorials/2026/07/mint-ai-illusion-of-wisdom-automation-2026/