International Journal of Contemporary Research In Multidisciplinary

Research Article

The Bedside in the Age of Artificial Intelligence: Why Pediatric Residency Still Needs Human Learning

Sabavath Arun [1*], Shivangi Sinha [2], Shivani Khandelwal [3], Shishir Shouri [4], Ajay Kumar [5]

[1,3,] Department of Paediatrics and Neonatology, All India Institute of Medical Sciences (AIIMS), Raebareli, Uttar Pradesh, India

[2,4,5] Department of Paediatrics, All India Institute of Medical Sciences (AIIMS), Bhopal Madhya Pradesh, India

Corresponding Author: * Sabavath Arun

DOI: https://doi.org/10.5281/zenodo.21186874

Abstract

Artificial intelligence (AI) is rapidly becoming embedded in the everyday practice of medicine and in the education of the physicians of tomorrow. Within seconds, a resident can now retrieve clinical guidelines, interrogate the latest evidence, generate structured differential diagnoses, and obtain decision support that once required hours in a library or the guidance of a senior colleague. These capabilities carry undeniable value, yet they have also intensified a longstanding anxiety about the future of traditional clinical training, and in particular about whether time-honoured methods such as bedside teaching retain their relevance. In paediatrics, this question carries special weight. Caring for children demands more than the recognition of disease patterns; it requires the ability to read a non-verbal patient, to earn the trust of frightened parents, to interpret subtle physical signs, and to make defensible decisions under uncertainty. Many of these competencies—clinical reasoning, communication, professionalism, and empathy—are acquired only through repeated, supervised, direct contact with patients and their families. AI can accelerate access to information and enrich the learning that surrounds a clinical encounter, but it cannot examine a breathless infant, comfort a distraught mother, or weigh the social realities that shape a treatment plan. This article argues that the future of pediatric residency lies not in choosing between technological innovation and bedside teaching, but in deliberately integrating the two so that each strengthens the other.

Manuscript Information

- **ISSN No:** 2583-7397
- **Received:** 11-05-2026
- **Accepted:** 29-06-2026
- **Published:** 04-07-2026
- **IJCRM:** 5(4); 2026: 36-40

- **Plagiarism Checked:** Yes
- **Peer Review Process:** Yes

How to Cite this Article

Arun S, Sinha S, Khandelwal S, Shouri S, Kumar A. The Bedside in the Age of Artificial Intelligence: Why Pediatric Residency Still Needs Human Learning. Int J Contemp Res Multidiscip. 2026;5(4):36-40.

Access this Article Online

www.multiarticlesjournal.com

KEYWORDS: Artificial Intelligence, Bedside Teaching, Pediatric Residency, Medical Education, Clinical Skills, Clinical Judgment, Empathy, Paediatrics.

1. INTRODUCTION

Over the past few years, artificial intelligence (AI) has moved from a speculative, futuristic concept into a routine presence at the point of care and in the classrooms and wards where physicians are trained. Tools that summarise the literature, propose differential diagnoses, interpret laboratory and imaging data, and draft clinical documentation are now available to clinicians and trainees at every level, often through nothing more than a smartphone [1]. Medical students and residents can review the most recent guidelines, appraise emerging evidence, and explore diagnostic possibilities with a speed and convenience that would have seemed extraordinary only a decade ago. As educators, we welcome innovations that genuinely improve learning and patient care, and there is much in the current wave of AI to welcome. Yet the accelerating presence of these tools has also revived an old and important question in a new form: as machines become more capable of retrieving and synthesising medical knowledge, will the slow, human, and resource-intensive practice of bedside teaching come to be seen as an inefficiency to be optimised away?

For pediatricians, this question cuts especially deep, because pediatric medicine has never been reducible to the recognition of a disease and the selection of a drug. Each working day brings encounters with children across a vast developmental range—from a silent neonate whose only language is physiology to an adolescent negotiating autonomy—alongside anxious parents and families confronting frightening and unfamiliar situations. A large share of the clinically decisive information in pediatrics is non-verbal and contextual: the quality of a cry, the effort of breathing, the way a toddler does or does not engage with a caregiver, the unspoken fear behind a parent's questions. These are lessons that cannot be found in a textbook or generated by an algorithm, and it is precisely because so much of pediatric assessment lives in this tacit, observational domain that bedside teaching continues to occupy a central place in pediatric training [1,2].

It is worth stating clearly at the outset that this article is not a defence of nostalgia, nor an argument against technology. The intention is neither to romanticise the traditional ward round nor to resist tools that demonstrably help residents learn and patients recover. Rather, our aim is to examine what bedside teaching actually does for a developing pediatrician—what specific competencies it builds and why those competencies matter—and then to ask honestly where AI can support that formation and where it cannot. Framing the debate as a contest between the bedside and the machine, in which one must eventually displace the other, misunderstands the nature of both. The more useful question is how a residency programme should combine them so that graduates emerge as clinicians who are simultaneously technologically fluent and deeply skilled in the human craft of caring for sick children.

In the sections that follow, we first consider the enduring educational value of the bedside: its role in transforming abstract knowledge into clinical understanding, in cultivating the pattern recognition that underlies the pediatric "sick or not sick" judgment, and in teaching the communication, professionalism, and empathy that define good pediatric care. We then turn to AI itself, acknowledging its genuine strengths as an educational and clinical aid while delineating its limits at the bedside. Finally, we propose a blended model in which technology is deliberately positioned to augment rather than replace direct patient contact, and we consider why this balance is particularly consequential in high-volume, resource-variable settings such as India [2,3].

2. DISCUSSION

Most paediatricians can recall a particular patient who taught them something that no lecture, textbook, or database ever could. These formative encounters are rarely dramatic; more often they are quiet moments of recognition that reshape how a trainee sees, thinks, and behaves. To understand why the bedside remains indispensable even in an era of powerful AI, it is helpful to separate out the distinct competencies it develops, because each stands in a different relationship to what technology can and cannot provide. We consider four in turn—the transformation of knowledge into understanding, the cultivation of clinical judgment, the acquisition of the human skills of medicine, and the appropriate role of AI—before drawing them together into a practical model for residency training.

From Knowledge to Clinical Understanding

A resident may read extensively about respiratory distress and be able to recite its causes, mechanisms, and management in impressive detail. Yet recognising the subtle chest retractions of an infant, hearing the difference between transmitted upper-airway sounds and true crepitations, or sensing that a child simply looks "sicker than the numbers suggest" is a capability that develops through seeing patients, not screens. The gap between propositional knowledge—knowing that—and practical knowledge—knowing how—is one of the oldest observations in medical education, and it is nowhere wider than in pediatrics, where the physical findings are often subtle, rapidly changing, and difficult to elicit from an uncooperative or frightened child. Bedside teaching is the setting in which this gap is closed. Under the guidance of an experienced clinician, the trainee learns to attach the words on the page to the reality in front of them: to feel the texture of a liver edge, to appreciate the tempo of a respiratory pattern, and to calibrate their own senses against those of a more experienced observer.

This transformation is gradual and cumulative. It depends on repeated exposure, on the opportunity to be corrected in real time, and on the presence of a teacher who can make explicit what would otherwise remain tacit—narrating why a particular finding matters, what it excludes, and how it should change the plan. As teachers, we have repeatedly observed residents gain genuine confidence only after they have examined patients themselves, even when they had already reviewed the relevant topic thoroughly beforehand. The information was not the missing ingredient; the embodied, supervised experience was. This is why the bedside remains the place where knowledge is metabolised into clinical understanding [2].

It is important to recognise that AI, however sophisticated, operates almost entirely in the domain of propositional knowledge. It can explain the pathophysiology of bronchiolitis with great fluency and can generate an exhaustive account of

the signs one might expect to find. What it cannot do is stand beside a trainee at the cot-side, watch the same infant, and say, in effect, "look here, feel this, notice how that changes when the child settles." The perceptual and psychomotor apprenticeship at the heart of clinical medicine has no substitute in text, and it is precisely this apprenticeship that bedside teaching provides. A residency that allowed information access to crowd out supervised patient contact would risk producing physicians who know a great deal and can do comparatively little.

Cultivating Clinical Judgment Under Uncertainty

Perhaps the single most important product of bedside learning is the development of clinical judgment. In real practice, patients rarely present exactly as they are described in textbooks. Symptoms are frequently incomplete or atypical, investigations may be equivocal or unavailable, and management decisions routinely have to be made before the picture is fully clear. The pediatric population magnifies each of these difficulties: a history is often obtained second-hand from a caregiver, the range of normal shifts with every month of age, and the same presenting complaint may signify a trivial viral illness in one child and an emerging emergency in another.

Learning to navigate this uncertainty is not a matter of memorising more facts; it is a matter of developing a disciplined way of thinking. At the bedside, and in the discussions that surround it, residents are pressed to articulate what they think is happening, to justify their reasoning, to state what would change their mind, and to defend a plan of action in the presence of ambiguity. A supervisor who asks "what else could this be, and how would you tell?" or "what are you most afraid of missing here?" is teaching a form of reasoning that generic information retrieval cannot supply. This is where the pediatric gestalt—the rapid, often pre-verbal judgment that a child is sick or not sick—is forged. It is a synthetic skill, built from thousands of small observations, and it frequently outperforms any single measured parameter in identifying the child who is deteriorating.

AI can genuinely support this process, but its role is auxiliary. A well-designed decision-support tool may broaden a differential, flag an overlooked possibility, or remind a tired resident of a guideline threshold. Yet the tool works from the data it is given, and in pediatrics the most important data are often those that a machine never receives: the child's colour and tone, the quality of parental concern, the trajectory over the last hour, and the countless contextual signals that an experienced clinician integrates almost unconsciously. Moreover, an over-reliance on algorithmic output carries its own hazard—the risk that a trainee accepts a plausible-sounding recommendation without developing the independent judgment needed to recognise when that recommendation is wrong. The capacity to disagree with a confident but mistaken source, human or machine, is itself a product of bedside experience. For this reason, discussing difficult and ambiguous cases at the bedside remains central to becoming a competent and safe pediatrician [3].

The Human Dimension: Communication, Professionalism, and Empathy

Beyond diagnosis and management, bedside teaching is where trainees learn the human side of medicine, which in pediatrics is inseparable from clinical competence itself. During ward rounds and clinic encounters, residents witness how parents react to uncertainty, how families cope—or struggle to cope—with serious illness, and how the words and manner of a clinician can either build or erode the trust on which treatment adherence ultimately depends. They learn that breaking difficult news to a family, negotiating a management plan with a reluctant parent, or comforting a frightened child are not soft extras appended to the "real" medicine but are integral to it. In pediatrics the patient and the decision-maker are usually different people, and caring for a child means caring for a whole family unit; this triangular relationship can only be learned by living inside it.

These human competencies are acquired largely through observation and supervised practice, in the same apprenticeship model that builds physical examination skills. A resident learns to sit rather than stand, to use language a parent can understand, to read the emotional temperature of a room, and to recognise when silence is more therapeutic than explanation, by watching skilled seniors do these things and then attempting them under gentle observation. Feedback on these moments—delivered privately, specifically, and with care—is one of the most valuable and least replaceable functions of a good bedside teacher. Over time, repeated exposure to families in distress, mediated by thoughtful supervision, is what shapes a technically capable trainee into a compassionate clinician.

It is precisely here that the limits of AI are most stark. A language model can draft a sympathetic-sounding paragraph, and can even suggest phrases for a difficult conversation, but it does not sit in the room, does not perceive the mother's rising fear, and does not bear the moral weight of the decision being made. It cannot model professionalism, cannot serve as a role model, and cannot form the relationship within which a trainee's own professional identity develops. Empathy in medicine is not merely the production of appropriate words; it is a disposition cultivated through real encounters with real suffering, and it is transmitted from clinician to trainee in a way that has no digital shortcut. A pediatric residency that neglected this dimension in favour of information efficiency would fail its graduates and, ultimately, its patients.

What AI Does Well — and Where It Stops

None of the foregoing should be read as a dismissal of AI, which offers real and growing value to residents and their teachers. Used well, AI tools are genuinely useful for reviewing the literature, checking and comparing guidelines, generating and structuring differential diagnoses, summarising a complex evidence base, and identifying recent developments that a busy trainee might otherwise miss. They can compress hours of searching into minutes, lower the barrier to evidence-based practice, and free time that can then be redirected towards patients. For residents working long hours under heavy service loads, these are not trivial benefits, and educators who ignore them do their trainees a disservice.

The strengths of AI, however, are concentrated in a particular part of the clinical task: the retrieval, organisation, and

synthesis of existing knowledge. Its outputs are generated from patterns in the data on which it was trained or the information it is given in a query, and they are only as good as those inputs. In pediatrics, this creates several important limitations. First, much of the decisive information is never captured in any dataset the model can see—the child's appearance, the subtleties of the examination, the social and family context—so the model reasons from a systematically incomplete picture. Second, AI systems can produce confident, fluent, and entirely incorrect output, and a trainee who has not developed independent judgment may be poorly placed to detect the error. Third, the evidence and training data underlying many tools are drawn disproportionately from high-income, adult, and Western populations, so their recommendations may fit imperfectly with the epidemiology, resource availability, and practice realities of a given setting.

There is also a subtler, developmental concern. Skills that are outsourced are skills that may never be acquired. If a trainee habitually reaches for an algorithm before forming their own impression, the internal cognitive work that builds diagnostic expertise may simply not happen. This is not an argument against using the tools; it is an argument for using them at the right moment and in the right way—after, rather than instead of, the trainee's own reasoning. The distinction matters because the goal of residency is not merely to arrive at correct answers today but to build a clinician capable of arriving at them independently for the next three decades.

In short, AI provides suggestions based on available data; it does not comfort a worried parent, notice a child's behaviour during an examination, appreciate the social circumstances that shape a treatment decision, or bear responsibility for the outcome. Recognising both what it does well and where it stops is the precondition for using it wisely in training.

Toward Integration: AI as Augmenter, Not Replacement

The most productive way to think about AI in pediatric residency is therefore not as a replacement for bedside teaching but as a means of strengthening it. The two are complementary because they occupy different parts of the learning cycle. Before rounds, a resident might use an AI platform to revise the physiology of a condition they are about to encounter, to prepare a focused set of questions, or to remind themselves of current guideline thresholds. After rounds, the same tools can be used to compare management options, to explore an unfamiliar diagnosis raised during the encounter, or to consolidate learning by testing their reasoning against the evidence. Positioned this way, AI does not compete with the patient encounter; it enriches the preparation and reflection that surround it.

The essential principle is one of sequence and subordination. The real learning still happens when knowledge is applied to an actual patient, under supervision, with all the messiness and uncertainty that real patients bring. Technology can assist clinical thinking, but it cannot substitute for clinical experience, and a training programme should be designed so that the tools serve the encounter rather than displacing it [2,3]. In practice this means protecting and prioritising direct patient contact and supervised bedside teaching time as the irreducible core of the curriculum, while treating AI as a powerful adjunct layered around that core.

Realising this vision has implications for how residents are taught and assessed. Trainees increasingly need explicit instruction in the appropriate and critical use of AI: how to frame a clinical question, how to appraise and verify an AI-generated answer, how to recognise the tool's blind spots, and how to preserve their own reasoning rather than surrendering it. This is a genuine competency in its own right, and one that faculty must be supported to teach, since many current supervisors trained in an era before these tools existed. Assessment, too, should continue to emphasise the abilities that only direct observation can capture—history taking, examination, communication, and reasoning at the bedside—so that the incentives of the programme reinforce rather than undermine its educational priorities.

Approached in this spirit, the arrival of AI need not diminish bedside teaching at all. On the contrary, by taking over some of the burden of information retrieval, it can potentially free clinical time and cognitive bandwidth for exactly the kind of supervised, human, hands-on learning that machines cannot provide. The danger lies not in the technology itself but in allowing convenience to quietly erode the harder, slower work on which competent practice depends.

Bedside Skills in High-Volume, Resource-Variable Settings

This balance carries particular weight in countries such as India, where the realities of clinical practice make strong bedside skills not merely desirable but essential. Patient loads are high, the spectrum of disease is broad, and the same ward may hold conditions that are common worldwide alongside others rarely seen in high-income settings. In many facilities, advanced imaging and laboratory support are limited, delayed, or unavailable, and diagnostic decisions continue to depend heavily on a careful history and a thorough physical examination. In such environments, the ability to elicit and interpret clinical signs is often the difference between a timely diagnosis and a missed one, and it cannot be delegated to a device.

The distribution and reliability of technology add a further consideration. AI tools depend on connectivity, on hardware, and on data infrastructure that are unevenly available across a large and diverse country. A paediatrician whose competence rests on solid bedside skills can practise safely in a district hospital with intermittent power and no decision-support software; one whose competence has been built around constant algorithmic assistance may be left exposed when that assistance is absent. Training that privileges bedside mastery therefore produces clinicians who are more robust and more readily transferable across the full range of settings in which paediatricians actually work.

There is also a question of fit. As noted, the evidence underlying many AI tools is drawn disproportionately from populations that differ epidemiologically and socially from those seen in much of India, and uncritical application of such tools risks recommendations that are subtly or substantially mismatched to local reality. A well-trained clinician, grounded in bedside experience and attuned to local disease patterns and

social circumstances, is best placed to use these tools intelligently—to take what is helpful and to recognise what does not apply. Far from making bedside skill less important, the arrival of AI in these settings arguably makes the clinician's independent judgment more important than ever, because someone must remain capable of deciding when the machine is right and when it is not [3].

CONCLUSION

Artificial intelligence is likely to become an ever more prominent part of pediatric training and practice, and this is, on balance, a development to be welcomed. Used thoughtfully, these tools can widen access to evidence, support clinical reasoning, reduce the drudgery of information retrieval, and give hard-pressed residents more time for the work that matters most. To resist them wholesale would be both futile and unwise. Yet the enthusiasm they rightly generate should not be allowed to obscure a simpler and more durable truth: the bedside remains the place where future paediatricians learn to observe carefully, to communicate with children and families, to think independently under uncertainty, and to care for their patients with genuine empathy.

These competencies—perceptual, cognitive, and moral—are acquired through direct, repeated, supervised contact with real patients, and they have no digital substitute. A resident learns to recognise a sick child, to comfort a frightened parent, to hold a difficult decision in the face of ambiguity, and to take responsibility for an outcome only by doing these things, again and again, under the guidance of experienced teachers. AI can prepare the ground and consolidate the lessons, but it cannot walk onto the ward in the trainee's place.

The goal, then, should never be framed as a choice between technology and bedside teaching, as though the ascent of one required the retreat of the other. It should be to use both wisely: to embrace AI as a powerful adjunct while fiercely protecting the supervised, hands-on, human learning that turns knowledge into competence and competence into care. Technology may help us reach a diagnosis faster, but it is the bedside that teaches us how to care for a child. Preserving that lesson—designing our residency programmes deliberately so that innovation strengthens rather than supplants it—should remain a central priority as pediatric education moves further into the age of artificial intelligence.

REFERENCES

1. Maske PG, Boddu LP, Hasamnis VA, Velidi MPS. Comparative study between bedside clinic versus simulation-based learning for medical students in the field of obstetrics and gynaecology. Eur J Cardiovasc Med. 2025;15:279-283. doi:10.5083/EJCM/25-06-53.
2. Roche E. Advances in undergraduate and postgraduate education: bedside, simulation and e-learning. Glob Pediatr. 2024;9:100188. doi:10.1016/j.gpeds.2024.100188.
3. Olaye IM, Seixas AA. The gap between AI and bedside: participatory workshop on the barriers to the integration, translation, and adoption of digital healthcare and AI startup technology into clinical practice. J Med Internet Res. 2023;25:e32962. doi:10.2196/32962.

About the Corresponding Author

Sabavath Arun is affiliated with the Department of Paediatrics and Neonatology, All India Institute of Medical Sciences (AIIMS), Raebareli, Uttar Pradesh, India. His academic interests include paediatrics, neonatology, medical education, and clinical research. He is actively involved in patient care, teaching, and research to improve child health outcomes.