

Artificial Intelligence in Acute Respiratory Distress Syndrome: Emerging Applications in Diagnosis, Risk Stratification, and Personalized Management

Dr. Tanmayee Siddharth Gajhas¹, Dr. Narendrasinh Solanki², Dr. Abhinav Lambe³, Dr. Pranav Deshmukh⁴

¹JR3 (Corresponding Author)

²Senior Consultant

³Associate Professor

⁴JR2

Abstract: ***Background:** Acute respiratory distress syndrome (ARDS) is a heterogeneous, life-threatening form of acute respiratory failure associated with substantial morbidity and mortality. Despite advances in lung-protective ventilation, prone positioning, and supportive critical care, early recognition, accurate phenotyping, individualized ventilatory management, and prognostic assessment remain challenging. Artificial intelligence (AI), particularly machine learning (ML) and deep learning, offers novel approaches for integrating complex clinical, physiological, imaging, waveform, and biomarker data. **Objective:** Our review examines current and emerging applications of AI in ARDS, with emphasis on diagnosis, early prediction, phenotyping, risk stratification, ventilation optimization, and prognosis. **Review:** AI-based models can analyze electronic health records, laboratory parameters, vital-sign trajectories, chest imaging, ventilator waveforms, and molecular biomarkers to identify patients at risk of ARDS before conventional clinical recognition. ML-based phenotyping may distinguish biologically and clinically meaningful subgroups, including hyperinflammatory and hypoinflammatory phenotypes. AI-assisted ventilation may facilitate individualized selection of tidal volume, positive end-expiratory pressure, inspired oxygen concentration, and detection of patient-ventilator asynchrony. Machine-learning models also demonstrate promising performance for mortality prediction and longitudinal risk assessment. However, limitations include retrospective development, dataset bias, inadequate external validation, limited interpretability, and uncertain clinical effectiveness. **Conclusion:** AI may transform ARDS management from syndrome-based supportive care toward dynamic, data-driven precision critical care. Prospective multicentre validation and clinically integrated trials remain essential before widespread implementation.*

Keywords: Artificial intelligence; Acute respiratory distress syndrome; Machine learning

1. Introduction

Acute respiratory distress syndrome (ARDS) is a severe form of acute hypoxaemic respiratory failure characterized by diffuse inflammatory lung injury, increased pulmonary vascular permeability, loss of aerated lung tissue, impaired gas exchange, and reduced respiratory-system compliance. The recently proposed global definition has broadened the conceptual framework of ARDS by accommodating high-flow nasal oxygen, pulse-oximetry-based oxygenation assessment under specified circumstances, and ultrasound as an imaging modality, thereby increasing applicability across diverse healthcare environments.[1] Nevertheless, ARDS remains a complex clinical syndrome rather than a single disease entity.

One major obstacle to improving ARDS outcomes is its marked heterogeneity. Patients satisfying the same diagnostic criteria may have different precipitating insults, inflammatory profiles, patterns of lung injury, recruitability, responses to positive end-expiratory pressure (PEEP), and clinical trajectories. Consequently, interventions effective for one biological or physiological subgroup may be ineffective or harmful in another.

Artificial intelligence (AI) offers an approach capable of processing this complexity. Machine learning (ML), deep learning, natural language processing, reinforcement

learning, and other computational techniques can analyze multidimensional data that are difficult to interpret simultaneously using conventional methods. Recent systematic reviews suggest increasingly promising applications of AI in ARDS diagnosis, prediction, subphenotyping, mortality assessment, and mechanical ventilation. [2-4]

Rather than replacing clinical judgment, AI may provide an additional layer of computational decision support. Its greatest potential lies in moving ARDS management from relatively static population-based protocols toward continuously updated, individualized risk assessment and treatment.

Artificial Intelligence in ARDS

AI is an umbrella term encompassing computational methods capable of performing tasks traditionally requiring human intelligence. ML represents a major component of AI in which algorithms identify relationships within data without requiring every decision rule to be explicitly programmed.

Common approaches used in critical care include logistic regression augmented by automated feature selection, random forests, gradient-boosting machines such as XGBoost, support-vector machines, artificial neural networks, convolutional neural networks (CNNs), recurrent neural networks, transformers, and reinforcement-learning systems.

ARDS is particularly suitable for AI-assisted management because modern intensive care generates enormous quantities of continuously changing data. These include heart rate, respiratory rate, oxygen saturation, blood pressure, arterial blood gases, laboratory results, ventilator settings, pressure and flow waveforms, imaging, medications, fluid balance, and increasingly genomic, proteomic and metabolomic information.

AI systems can combine these data streams and potentially identify patterns that precede clinically obvious deterioration.

AI in Early Prediction of ARDS

Early recognition is crucial because pathological lung injury may already be progressing before patients fulfil formal diagnostic criteria. Conventional recognition relies on clinical suspicion, oxygenation indices, imaging findings, and identification of an appropriate precipitating insult. AI provides an opportunity to identify high-risk patients earlier.

Models developed using electronic health record (EHR) data typically incorporate demographics, vital signs, oxygen requirements, arterial blood gases, laboratory parameters, comorbidities, organ dysfunction, mechanical ventilation variables, and temporal changes in these measurements.

A 2025 systematic review and meta-analysis evaluating ML models for ARDS prediction reported a pooled area under the receiver-operating-characteristic curve of approximately 0.74, with substantial between-study heterogeneity.[5] These findings demonstrate useful discriminatory potential while simultaneously emphasizing that apparently high performance in individual derivation studies may not translate uniformly across institutions.

The major advantage of AI is its capacity for **dynamic prediction**. Instead of estimating risk only at ICU admission, algorithms can continuously recalculate the probability of ARDS as physiological and laboratory data accumulate.

For example, progressive oxygen requirement, increasing respiratory rate, declining PaO₂/FiO₂ ratio, inflammatory abnormalities, haemodynamic instability, and changes in ventilatory requirements could collectively generate an early-warning signal before overt deterioration.

Such prediction could permit earlier clinical reassessment, closer monitoring, evaluation of potentially reversible causes, and consideration of lung-protective strategies in patients likely to require invasive ventilation.

AI prediction, however, should not be interpreted as an autonomous diagnostic system. False-positive alerts may contribute to alarm fatigue, whereas false-negative predictions could provide inappropriate reassurance. External validation, calibration, and demonstration of actual clinical benefit are therefore essential.

AI-Assisted Diagnosis

ARDS diagnosis remains partly dependent on clinician interpretation. Radiographic infiltrates may be difficult to distinguish from cardiogenic pulmonary oedema, atelectasis, pneumonia, or other causes of respiratory failure.

Deep-learning methods, particularly CNNs and foundation models, can extract imaging features from chest radiographs and computed tomography (CT) that may be difficult to quantify visually.

A recent systematic review and meta-analysis of AI models for ARDS diagnosis and subphenotyping reported pooled sensitivity of approximately 0.89, specificity of 0.88, and an AUROC of approximately 0.90, although heterogeneity was high and external validation was frequently absent.[6]

The potential role of imaging AI is rapidly expanding. Automated systems can quantify normally aerated, poorly aerated and non-aerated lung regions and characterize the spatial distribution of pulmonary injury. Such quantitative imaging may ultimately complement conventional oxygenation criteria.

A particularly recent development is CT-based integrated AI. AutoARDS, reported in 2026, was developed using more than 50,000 CT volumes and validated across six centres involving more than 6,000 individuals. The model was designed not simply to recognize ARDS but also to derive quantitative information regarding disease severity, oxygenation and prognosis from CT imaging.[7] Such systems illustrate how imaging could evolve from qualitative pattern recognition into a quantitative physiological biomarker.

Artificial Intelligence and ARDS Phenotyping

Why Phenotyping Matters

Perhaps the most important precision-medicine application of AI in ARDS is **phenotyping**.

ARDS represents a final common clinical syndrome arising from multiple biological mechanisms. Treating all patients as a homogeneous population may partially explain why numerous pharmacological trials have failed despite promising biological rationale.

Latent class analysis and related computational techniques have repeatedly identified two broadly reproducible ARDS phenotypes: a **hyperinflammatory phenotype** and a **hypoinflammatory phenotype**. [8]

The hyperinflammatory phenotype generally demonstrates greater systemic inflammation, haemodynamic instability and organ dysfunction and is associated with poorer outcomes. The hypoinflammatory phenotype generally has lower inflammatory biomarker concentrations and more favourable outcomes.

Importantly, secondary analyses of previous clinical trials have suggested that these phenotypes may demonstrate heterogeneous responses to interventions. Thus, phenotype identification could eventually become relevant not only for prognosis but also for treatment selection.

AI-Based Phenotype Identification

Traditional phenotype identification may require biomarkers that are not immediately available at the bedside. ML offers the possibility of identifying comparable subgroups using routinely collected clinical information.

Sinha and colleagues demonstrated that ML classifier models could identify previously described ARDS phenotypes using readily available clinical variables, potentially reducing dependence on specialized biomarker measurements.[9]

More recent evidence continues to support computational phenotyping. A 2026 scoping review found that the distinction between hyperinflammatory and hypoinflammatory ARDS remains the most consistently reproduced phenotypic framework and is associated with differences in mortality, ventilator-free days and organ failure.[10]

AI also allows phenotype classification to become **dynamic rather than static**. A patient's biological and physiological state can change during ICU treatment. Recent longitudinal subphenotyping research has demonstrated distinct temporal trajectories and transitions between ARDS states, with transitions themselves carrying prognostic information.[11]

This concept is highly relevant. Future precision treatment may depend not simply on asking, "Which phenotype does this patient have?" but rather, "Which phenotype does this patient have now, and how is that phenotype evolving?"

AI for Risk Stratification

Risk stratification traditionally uses indices such as APACHE, SOFA and SAPS, together with ARDS severity based largely on oxygenation. These systems are clinically useful but cannot fully capture nonlinear relationships among hundreds of physiological variables.

ML algorithms can incorporate much larger numbers of variables and their interactions. Dynamic information such as trends in oxygenation, airway pressure, acid-base status, inflammatory parameters and organ dysfunction may be more informative than isolated measurements.

Risk stratification has several potential clinical applications:

- 1) Identifying patients likely to deteriorate;
- 2) Prioritizing intensive monitoring;
- 3) Identifying candidates for advanced respiratory support;
- 4) Informing discussions regarding prognosis;
- 5) Enriching clinical trials with patients most likely to experience relevant outcomes; and
- 6) Identifying subgroups likely to respond differently to treatment.

AI therefore enables both **prognostic enrichment**, in which patients are selected according to outcome risk, and **predictive enrichment**, in which patients are selected according to their likelihood of responding to a particular intervention.

AI in Mechanical Ventilation Optimization

Mechanical ventilation is lifesaving in severe ARDS but can itself worsen pulmonary injury through excessive tidal volumes, pressures, cyclic opening and closing of unstable lung units, and excessive mechanical power.

Current lung-protective ventilation provides population-level principles, but optimal settings may differ considerably between individual patients.

AI has considerable potential to facilitate personalized ventilation by integrating respiratory-system mechanics, blood gases, oxygen saturation, haemodynamic variables, ventilator waveforms and previous responses to ventilator adjustments.

Tidal Volume and PEEP

Algorithms can potentially predict the physiological consequences of alternative combinations of tidal volume, respiratory rate, FiO₂ and PEEP.

Rather than selecting PEEP from a fixed table alone, future systems could integrate compliance, driving pressure, oxygenation response, haemodynamics, recruitability and imaging characteristics to estimate individualized settings.

Reinforcement learning is particularly interesting because it learns sequential treatment strategies based on states, actions and subsequent outcomes. A 2024 retrospective study using eICU and MIMIC-IV data developed a reinforcement-learning strategy recommending different levels of PEEP, FiO₂ and ideal-body-weight-adjusted tidal volume according to patient condition.[12] The findings suggested potential advantages over observed practice, although prospective clinical trials are required before such algorithms can be used autonomously.

Thus, current evidence should be interpreted as proof of concept rather than evidence that AI-selected ventilator settings improve survival.

Detection of Patient–Ventilator Asynchrony

Patient–ventilator asynchrony occurs when the patient's neural respiratory effort is poorly matched with ventilator-delivered support. Double triggering, ineffective efforts, auto-triggering, premature cycling and delayed cycling may increase discomfort and potentially contribute to injurious ventilation.

Manual recognition requires expertise and continuous waveform observation.

AI can continuously analyze pressure, flow and volume waveforms and automatically identify abnormal breath patterns.

A recent review of AI for patient–ventilator asynchrony identified 13 studies encompassing more than 5.8 million breaths. Many models demonstrated sensitivity and specificity exceeding 90%, highlighting the potential of automated waveform analysis.[13]

Continuous AI monitoring could therefore detect clinically important asynchronies that might otherwise remain unnoticed between bedside assessments.

Mechanical Power and Lung Injury Prevention

Mechanical power integrates several determinants of ventilator-associated energy delivery, including tidal volume, respiratory rate, airway pressure, PEEP and flow. Excessive energy delivered to vulnerable lung tissue may contribute to ventilator-induced lung injury.

AI systems could continuously calculate mechanical-power-related variables and identify combinations of ventilator settings associated with potentially injurious energy exposure.

Future closed-loop systems may simultaneously optimize oxygenation, carbon dioxide clearance, driving pressure, mechanical power and haemodynamic tolerance rather than adjusting each variable independently.

AI in Prognosis and Mortality Prediction

Accurate prognosis in ARDS is difficult because outcomes depend on pulmonary injury, systemic illness, comorbidities, age, organ failure and treatment response.

ML models can integrate these factors to generate individualized mortality estimates.

A 2025 systematic review and meta-analysis involving 21 studies and more than 31,000 patients found that ML-based models showed strong discrimination for ARDS mortality. The pooled C-index in external validation datasets was approximately 0.81, while traditional scoring systems demonstrated lower overall discrimination.[14]

Dynamic prediction appears particularly valuable. In an analysis of ARDS Network FACTT data, an ML model using day-3 clinical information achieved better mortality discrimination than a model based on baseline information, illustrating how evolving physiological response contains important prognostic information.[15]

Variables contributing to prognosis include age, oxygenation, respiratory mechanics, airway pressures, haemodynamic parameters, platelet count, albumin, renal function, inflammatory markers, acid-base status and the extent of extrapulmonary organ dysfunction.

AI could eventually generate continuously updated risk trajectories rather than a single mortality probability at admission.

Such estimates should nevertheless support- not replace- clinical judgment. Prognostic algorithms can be affected by treatment-selection bias and historical patterns of care. Consequently, predictions must not become self-fulfilling, particularly when decisions regarding escalation or limitation of life-sustaining therapy are being considered.

Multimodal AI and Personalized ARDS Management

The next generation of ARDS AI is likely to be multimodal.

A single patient may generate:

- Structured EHR data,
- Laboratory measurements,
- Arterial blood gases,
- Radiographic and CT images,
- Ventilator waveforms,
- ECG signals,
- Bedside ultrasound,
- Clinical notes,
- Inflammatory biomarkers, and
- Genomic, transcriptomic, proteomic or metabolomic information.

Multimodal AI could integrate these heterogeneous sources into a unified representation of the patient's physiological state.

This creates the possibility of a continuously updated **digital phenotype** of ARDS.

For example, an AI platform could identify increasing ARDS probability, classify inflammatory phenotype, quantify CT injury, evaluate recruitability, recognize ventilator asynchrony, estimate mortality risk and recommend consideration of prone positioning or reassessment of ventilator settings.

Such an approach represents a transition from isolated prediction models toward comprehensive clinical decision-support systems.

Large Language Models and Generative AI

Large language models (LLMs) represent another emerging AI application. They can process unstructured information contained in progress notes, radiology reports, nursing documentation and clinical narratives.

In ARDS care, LLMs could potentially summarize longitudinal patient information, identify missing diagnostic criteria, extract risk factors from clinical notes, explain model outputs, and present complex multimodal information in clinician-friendly formats.[2]

However, generative AI can produce inaccurate or fabricated information. Consequently, LLM-generated recommendations require explicit source grounding, clinical oversight and rigorous validation before integration into high-risk critical-care decision making.

2. Challenges and Limitations

Despite rapid progress, several barriers prevent routine clinical implementation.

- 1) **Data quality and missingness:** ICU datasets contain measurement errors, missing observations and inconsistent documentation.
- 2) **Generalizability:** Models trained in one hospital or database may perform poorly in different populations, ventilator systems or healthcare environments.
- 3) **Algorithmic bias:** Under-represented populations may experience poorer model performance.
- 4) **Interpretability:** Complex deep-learning models may generate accurate predictions without providing clinically understandable explanations.
- 5) **Calibration:** A model may discriminate survivors from non-survivors while producing inaccurate absolute probabilities.
- 6) **Dataset shift:** Changes in clinical practice, diagnostic criteria and treatment protocols can reduce model performance over time.
- 7) **External validation:** Many published models have been internally validated but lack robust independent multicentre validation.
- 8) **Clinical utility:** High AUROC does not necessarily translate into improved patient outcomes. A clinically

useful AI tool must alter decisions appropriately and improve meaningful outcomes.

- 9) **Automation bias:** Clinicians may place excessive confidence in algorithmic recommendations.
- 10) **Ethics and privacy:** Continuous use of detailed clinical data raises concerns regarding consent, cybersecurity, data ownership and governance.

Recent evaluation of AI-driven decision support in acute respiratory failure found surprisingly few studies assessing real-world clinical outcomes, despite the large number of model-development studies.[16] This translational gap remains one of the most important limitations of the field.

3. Conclusion

Artificial intelligence has emerged as a promising tool for addressing several persistent challenges in ARDS. Its applications extend from early prediction and automated diagnosis to phenotyping, risk stratification, mechanical ventilation optimization and prognostic assessment.

Among these applications, AI-enabled phenotyping may be particularly important because it addresses the fundamental biological heterogeneity of ARDS. Dynamic ML models may identify hyperinflammatory, hypoinflammatory and other clinically meaningful phenotypes and eventually help match therapies to individual biological and physiological characteristics.

References

- [1] Matthay MA, Arabi Y, Arroliga AC, Bernard G, Bersten AD, Brochard LJ, et al. A new global definition of acute respiratory distress syndrome. **Am J Respir Crit Care Med.** 2024;209(1):37-47. doi:10.1164/rccm.202303-0558WS.
- [2] Wang Z, et al. Artificial intelligence and machine learning in acute respiratory distress syndrome management: recent advances. **Front Med (Lausanne).** 2025; 12: 1597556. doi:10.3389/fmed.2025.1597556.
- [3] Suriyaarachchi P, et al. Machine learning in ARDS: an intensivist's guide to artificial intelligence applications. **[Journal details indexed in PubMed].** 2026.
- [4] Choudhury A, et al. Artificial intelligence in acute respiratory distress syndrome: a systematic review. **Int J Environ Res Public Health.** 2022;19. PMID:36100348.
- [5] Yang J, Zeng S, Cui S, Zheng J, Wang H. Predictive modeling of acute respiratory distress syndrome using machine learning: systematic review and meta-analysis. **J Med Internet Res.** 2025; 27: e66615. doi:10.2196/66615.
- [6] Rao PK, Systematic review and meta-analysis of artificial intelligence models for diagnosing and subphenotyping ARDS in adults. **[Journal indexed in PubMed].** 2025. PMID:41037977.
- [7] Wans BB, CT-based AI system for quantitative and integrated management of acute respiratory distress syndrome in critical care. **2026.** PMID:42032151.
- [8] Sinha P, Calfee CS. Phenotypes in acute respiratory distress syndrome: moving towards precision medicine. **Curr Opin Crit Care.** 2019;25(1):12-20. doi:10.1097/MCC.0000000000000571.
- [9] Sinha P, Churpek MM, Calfee CS. Machine learning classifier models can identify acute respiratory distress syndrome phenotypes using readily available clinical data. **Am J Respir Crit Care Med.** 2020;202(7):996-1004. doi:10.1164/rccm.202002-0347OC.
- [10] Johnson, Subphenotypes in acute respiratory distress syndrome: a scoping review across clinical, biological, computational, imaging, omics, and artificial intelligence approaches. **2026.** PMID:41619711.
- [11] Komali P, Longitudinal subphenotyping of acute respiratory distress syndrome: implications in prognostic prediction and clinical intervention. **2026.** PMID:41865760.
- [12] Digjam, Reinforcement learning to optimize ventilator settings for patients on invasive mechanical ventilation: retrospective study. **J Med Internet Res.** 2024; 26: e44494. doi:10.2196/44494.
- [13] Hatson K, Artificial intelligence for the detection of patient-ventilator asynchrony. **Respir Care.** 2025. doi:10.1089/respcare.12540.
- [14] Tan R, Ge C, Li Z, Yan Y, Guo H, Song W, et al. Early prediction of mortality risk in acute respiratory distress syndrome: systematic review and meta-analysis. **J Med Internet Res.** 2025. PMID:40392588.
- [15] Kunal BB, Early predictive values of clinical assessments for ARDS mortality: a machine-learning approach. **Sci Rep.** 2024. doi:10.1038/s41598-024-68653-8.
- [16] PAUL B, Artificial intelligence-driven decision support for patients with acute respiratory failure: a scoping review. **Intensive Care Med Exp.** 2025;13(1):83. doi:10.1186/s40635-025-00791-3.