WHAT'S NEW IN INTENSIVE CARE

The future of artificial intelligence in intensive care: moving from predictive to actionable AI

Jim M. Smit[1,2]* , Jesse H. Krijthe[2] and Jasper van Bommel[1] on behalf of the Causal Inference for ICU Collaborators

Artificial intelligence (AI) research in the intensive care unit (ICU) mainly focuses on developing models (from linear regression to deep learning) to predict outcomes, such as mortality or sepsis [1, 2]. However, there is another important aspect of AI that is typically not framed as AI (although it may be more worthy of the name), which is the prediction of patient outcomes or events that would result from different actions, known as causal inference [3, 4]. This aspect of AI is crucial for decision-making in the ICU. To emphasize the importance of causal inference, we propose to refer to any data-driven model used for causal inference tasks as 'actionable AI', as opposed to 'predictive AI', and discuss how these models could provide meaningful decision support in the ICU.

Predictive versus actionable AI

Predictive AI should perform prediction tasks [3]. In the context of clinical practice, this involves generating forecasts of how likely patient outcomes are now, or in the future. As such, predictive AI can offer an early warning of possible adverse events, enabling ICU physicians to consider appropriate interventions pre-emptively. What it cannot do, is forecast how the probability of a patient's outcome might change if a particular intervention is implemented, as it relies entirely on associations [5]. For instance, while palliative care consults and norepinephrine infusions are both highly indicative of patient mortality, it is not reasonable to conclude that discontinuing either treatment would decrease the patient's probability of death [6]. In other words, predictive AI cannot guide ICU clinicians in what to do, as it solely offers an early warning.

For an AI to advise ICU physicians in treatment decisions, i.e., 'actionable AI', cause and effect need to be taken into account. Actionable AI should perform causal inference tasks [3], which means that it predicts (future) patient outcomes or events that would result from alternative treatment decisions. By comparing these outcomes, an actionable AI could advise on treatment options that lead to the best predicted outcome (i.e., the optimal treatment). In medicine, causal inference tasks are traditionally performed by conducting randomized controlled trials (RCTs). The randomization of the treatment allows one to interpret the difference in outcome between treatment arms as a causal effect of the treatment. Hence, one can simply compare outcomes and conclude that the one with the best observed outcome represents the optimal treatment. However, in observational studies, causal inference tasks are more complex, often compounded by bias stemming from common causes (confounding bias) and selection on common effects (selection bias). Thus, for an AI to 'learn' causal inference tasks from observational data, it needs to adjust for these biases. To do so, it is key to use an adjustment method that suits the type of treatment being considered.

*Correspondence: j.smit@erasmusmc.nl
[1] Department of Intensive Care, Erasmus University Medical Center, Doctor Molewaterplein 40, (internal postal address: Room Ne-411), 3015 GD Rotterdam, The Netherlands
Full author information is available at the end of the article

Springer

Intensive care medicine is about sequential decision-making

If a treatment decision only occurs at baseline, e.g., the randomized 'intention-to-treat' in an RCT, it is called a 'time-fixed' (or 'point') treatment. In observational studies considering time-fixed treatments, the treatment decision occurs only at baseline and, consequently, confounding occurs only at baseline as well. Here, conventional bias adjustment methods (like regression or propensity matching) suffice. However, ICU treatments typically comprise sequences of treatment decisions. In sepsis, for instance, the decision whether or not to administer fluids and vasopressors needs to be made not only at sepsis onset, but at multiple time-points during the ICU stay (Fig. 1). Hence, ICU patients are typically treated according to a certain regime—or policy, or strategy, or bundle—which represents a set of rules informing treatment decisions during follow-up, based on a patient's response. For example, liberal and restrictive fluid therapy represent two different regimes that dictate administration of fluids and vasopressors during sepsis onset. When there are multiple decisions through time, patient characteristics that act as confounders may vary over time and even be affected by previous treatment decisions, leading to the so-called 'time-varying confounding' [7]. To adjust for this appropriately, more sophisticated methods are required [8], some of which have been applied to ICU topics. We will discuss some examples and remaining challenges.

Actionable AI in the ICU: are we there yet?

Despite the predominant focus on predictive AI in ICU research, there is growing interest in developing actionable AI. For example, Shahn and colleagues [9] performed a 'target trial emulation' [10] to develop a marginal structural model that suggests sepsis outcomes could improve through more restricted fluid strategies. Similarly, Komorowski and colleagues [11] have presented a reinforcement learning model that predicts the optimal dosing of fluids and vasopressors in sepsis. Both models aim to perform causal inference tasks, despite Shahn's statistical approach and Komorowski's machine learning (ML) approach. However, both studies rely on observational data, and therefore, neither statistical nor ML methods can guarantee successful causal inference. Causal inference using observational data is challenging, and clinical domain knowledge is essential to understanding the causal relationships between treatment and outcome. Causal diagrams [12] can help visualize potential sources of bias, but it is crucial to note that bias can never be entirely ruled out using observational data. Moreover,
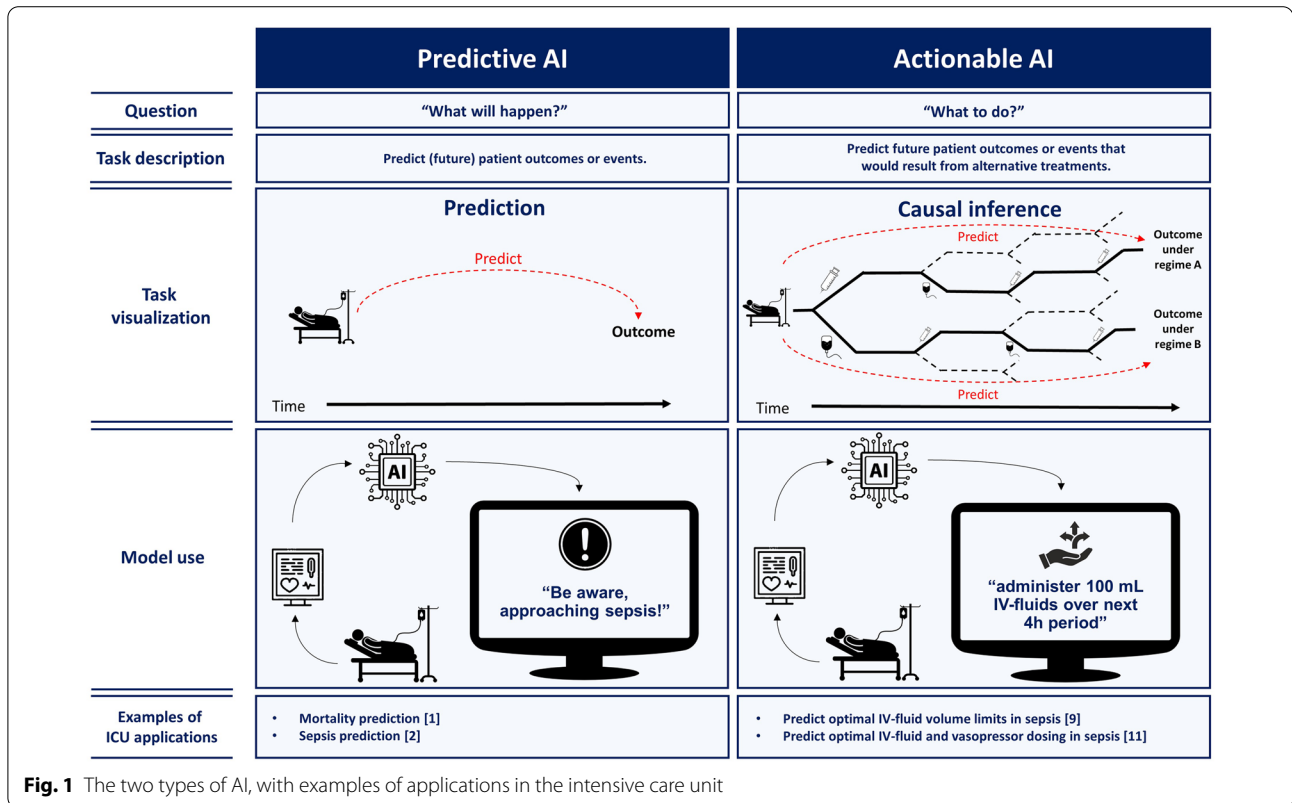

	Predictive AI	**Actionable AI**
Question	"What will happen?"	"What to do?"
Task description	Predict (future) patient outcomes or events.	Predict future patient outcomes or events that would result from alternative treatments.
Task visualization	Prediction	Causal inference
Model use	"Be aware, approaching sepsis!"	"administer 100 mL IV-fluids over next 4h period"
Examples of ICU applications	• Mortality prediction [1] • Sepsis prediction [2]	• Predict optimal IV-fluid volume limits in sepsis [9] • Predict optimal IV-fluid and vasopressor dosing in sepsis [11]

Fig. 1 The two types of AI, with examples of applications in the intensive care unit

a significant challenge is the typically limited 'effective sample size', which refers number of patient histories for which the modeled and observed treatment regimes agree [13]. Surmounting these challenges is a prerequisite for successfully implementing actionable AI in clinical practice. In our recent systematic review, we offer recommendations for future causal inference research using observational data in the ICU [14].

While the use of observational data is promising, towards actionable AI at the bedside, usage of RCT data may currently be the safest route. This is because in RCTs, the task of inferring causation is already achieved through randomization. Although RCTs offer estimates of average treatment effects, one can utilize these data to create models that produce more individualized estimates of treatment effects, i.e., 'personalized' or 'precision' medicine. However, even with enough RCT data available, the appropriate modeling approach is not straightforward, and still, various challenges remain [15].

Future perspective

Actionable AI models have the potential to guide ICU physicians in treatment choices, although challenges remain before these can safely be implemented. Rather than an omnipotent 'general AI' overtaking the control of the entirety of clinical decisions in the ICU, we envision that future actionable AI models will remain examples of 'narrow AI', confined at advising specific treatment decisions for specific patient groups and clinical scenarios.

Author details
[1] Department of Intensive Care, Erasmus University Medical Center, Doctor Molewaterplein 40, (internal postal address: Room Ne-411), 3015 GD Rotterdam, The Netherlands. [2] Pattern Recognition and Bio-Informatics Group, EEMCS, Delft University of Technology, Delft, The Netherlands.

Declarations

Conflicts of interest
On behalf of all authors, the corresponding author states that there is no conflict of interest.

References
1. van de Sande D, van Genderen ME, Huiskens J et al (2021) Moving from bytes to bedside: a systematic review on the use of artificial intelligence in the intensive care unit. Int Care Med 47:750–760. https://doi.org/10.1007/s00134-021-06446-7
2. Fleuren LM, Klausch TLT, Zwager CL et al (2020) Machine learning for the prediction of sepsis: a systematic review and meta-analysis of diagnostic test accuracy. Int Care Med 46:383–400. https://doi.org/10.1007/s00134-019-05872-y
3. Hernán MA, Hsu J, Healy B (2019) A second chance to get causal inference right: a classification of data science tasks. Chance 32:42–49. https://doi.org/10.1080/09332480.2019.1579578
4. Prosperi M, Guo Y, Sperrin M, et al. Causal inference and counterfactual prediction in machine learning for actionable healthcare. Nat Mach … Published Online First: 2020.https://www.nature.com/articles/s42256-020-0197-y
5. Savage N (2023) Why artificial intelligence needs to understand consequences. Nature. https://doi.org/10.1038/d41586-023-00577-1
6. Chen JH, Asch SM (2017) Machine learning and prediction in medicine — beyond the peak of inflated expectations. N Engl J Med 376:2507–2509. https://doi.org/10.1056/nejmp1702071
7. Mansournia MA, Etminan M, Danaei G et al (2017) Handling time varying confounding in observational research. BMJ 359:1–6. https://doi.org/10.1136/bmj.j4587
8. Daniel RM, Cousens SN, De Stavola BL et al (2013) Methods for dealing with time-dependent confounding. Stat Med 32:1584–1618. https://doi.org/10.1002/sim.5686
9. Shahn Z, Shapiro NI, Tyler PD et al (2020) Fluid-limiting treatment strategies among sepsis patients in the ICU: A retrospective causal analysis. Crit Care. https://doi.org/10.1186/s13054-020-2767-0
10. Hernán MA, Robins JM (2016) Using big data to emulate a target trial when a randomized trial is not available. Am J Epidemiol 183:758–764. https://doi.org/10.1093/aje/kwv254
11. Komorowski M, Celi LA, Badawi O et al (2018) The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care. Nat Med 24:1716–1720. https://doi.org/10.1038/s41591-018-0213-5
12. Tennant PWG, Murray EJ, Arnold KF et al (2021) Use of directed acyclic graphs (DAGs) to identify confounders in applied health research: review and recommendations. Int J Epidemiol 50:620–632. https://doi.org/10.1093/ije/dyaa213
13. Gottesman O, Johansson F, Komorowski M et al (2019) Guidelines for reinforcement learning in healthcare. Nat Med 25:16–18. https://doi.org/10.1038/s41591-018-0310-5
14. Smit JM, Krijthe JH, van Bommel J et al (2022) Causal inference using observational intensive care unit data: a systematic review and recommendations for future practice. Medrxiv. https://doi.org/10.1101/2022.10.29.22281684
15. Kent DM, Paulus JK, Van Klaveren D et al (2020) The predictive approaches to treatment effect heterogeneity (path) statement. Ann Intern Med 172:35–45. https://doi.org/10.7326/M18-3667