Intensive Care Med (2019) 45:1298–1300
https://doi.org/10.1007/s00134-019-05662-6

WHAT'S NEW IN INTENSIVE CARE

Artificial intelligence in intensive care: are we there yet?

Matthieu Komorowski*

Artificial intelligence (AI) as a novel tool to advance the field of medicine has rapidly become a subject of great interest and intense debate, leading many stakeholders to prospect on the future role of these technologies and relative responsibilities in decision-making [1–5]. Furthermore, the complexity of some AI algorithms, their lack of transparency and a widespread lack of prospective validation may serve to dishearten physicians. In this short piece, I will attempt to define what are some of the most common AI techniques with applications to intensive care, then I'll discuss what barriers to implementation exist and why us, doctors, have not been superseded by robot-physicians.

Intensive care applications of artificial intelligence

The computer scientist Marvin Minsky, one of the fathers of AI, defined the field as "the science of making machines do things that would require intelligence if done by men". AI has many potential applications in medicine, including automated radiology reporting, predictive analytics, knowledge representation, surgical robotics, medical education, drug discovery, among others [1–5]. These technologies hold the promises of helping clinicians make better and more personalised decisions, boosting workflow and reducing healthcare expenditures [1, 3–5]. In the ICU, applications of AI are mainly limited to machine learning—a collection of data analysis and modelling techniques aiming at generating knowledge from data. Machine learning algorithms typically fall into three categories: supervised, unsupervised and reinforcement learning (Fig. 1).

Supervised learning relies on labelled data for model training and is focused on predictive tasks such as estimating a patient's risk of mortality, readmission, length of stay, or diagnosing early deterioration or sepsis [1, 6, 7]. A vast overlap exists between biostatistics and machine learning, so one could argue that we have actually been using some form of AI for a long time. For example, many risk prediction scores (PRISM, APACHE, etc.) are all supervised models that are already used in clinical practice.

Figure 1 includes indications on the technology readiness level (TRL) of the methods. TRL is a scale of the maturity of technologies, originally invented by NASA but now widely adopted by many institutions as innovation policy tools [8]. The TRL of most supervised learning models published in intensive care medicine ranges from 4 (a component validated in laboratory) to 7 (mature system adequately validated in the real world), even though the majority of the published models have never been tested or deployed in clinical practice, what we could refer to as an "implementation bottleneck".

Next, unsupervised learning enables the discovery of latent structure or subclasses in a dataset. In medicine, this appears useful to define patient subgroups and phenotypes, which may be distinct from traditional subgroups [9]. These tools may also help making sense of "omics" data [1, 10]. For instance, Antcliffe analysed transcriptomics to define two "sepsis response signatures" associated with different steroid response [10]. While the potential of these technologies in "precision medicine" or to inform clinical trial design [2, 11] is vast, no prospectively validated application exists to this day, to the best of my knowledge. Arguably, most of these tools are TRL 6 at best ("System adequacy validated in simulated environment").

The final class of machine learning algorithms is reinforcement learning, where a virtual agent is trained to learn an optimal set of rules in order to maximise some reward [12, 13]. Researchers have made a parallel with

*Correspondence: m.komorowski14@imperial.ac.uk
Department of Surgery and Cancer, Faculty of Medicine, Imperial College London, London SW7 2AZ, UK

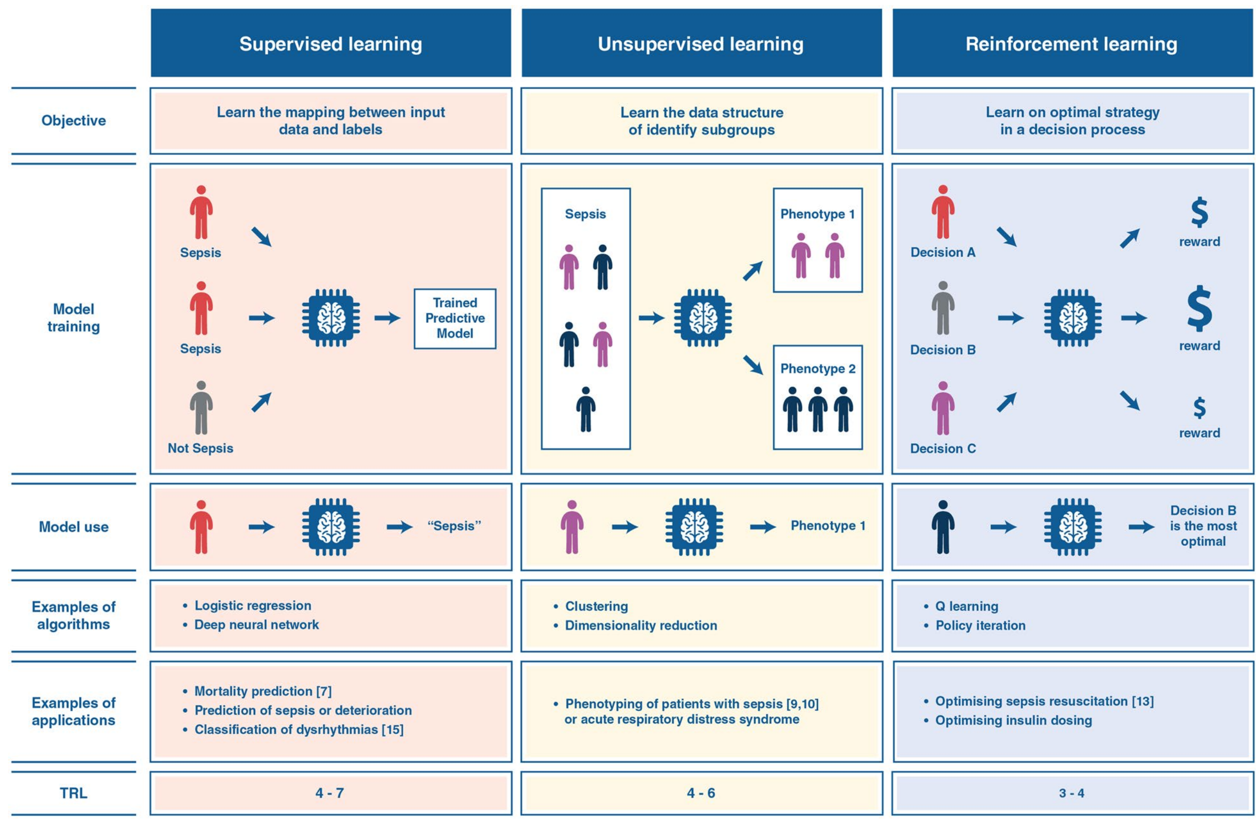

Fig. 1 The three classes of machine learning algorithms, with examples of applications in intensive care. Technology readiness level (TRL) refers to the maturity of a technology [8]. TRL 3: proof-of-concept demonstrated analytically and/or experimentally; TRL 4: component validated in the laboratory; TRL 6: system adequacy validated in simulated environment; TRL 7: mature system adequately validated in the real world

medicine and attempted to deploy these algorithms to solve clinical questions [13, 14]. This branch of machine learning is arguably the most immature, so the TRL of these algorithms remains low, of 3 ("proof-of-concept demonstrated analytically and/or experimentally") to 4 (laboratory-validated). More examples of applications of machine learning are referenced in the Supplementary Material.

Barriers to implementation

While AI has already deeply transformed many industries (retail, transportation, finance, etc.), healthcare is still lagging behind for a number of reasons [1–3]. First, the technological challenges are huge: setting up large databases, dealing with issues of interoperability and incomplete, spurious and artefacted data in order to build these complex models and turning them into bedside tools is inherently very difficult [2, 3]. Next, proprietary, legal, ethical and privacy issues need to be addressed [2]. Medicine is a high-risk environment, where deploying inaccurate or poorly calibrated algorithms would be unacceptable [13]. Institutions are trying to set up the correct regulatory framework, for example in the UK with the governmental report "Algorithms in decision-making" [15]. Finally, the successful deployment of such tools requires the trust and buy-in of all stakeholders [1, 2].

Clinicians will be mostly interested in scrutinising the explainability of a model (can the working of the algorithm be understood by humans?) and its generalisability (is the model usable in a given population?) [1, 13]. Explainability has been enforced for algorithms with human impact by the European General Data Protection Regulation (GDPR) since May 2018. The debate of whether it is acceptable to use non-transparent algorithms for patient care is ongoing [1]. Regarding generalisability, the same principle as any clinical trial applies to machine learning: a model is only applicable to the population on which it was built, and external validation (in another dataset or cohort, ideally from a different organisation or country) enhances greatly the robustness of a proposed method.

It is, however, reassuring to observe that these hurdles have begun to fall, and that innovations are starting to trickle down into clinical practice. In fact, some medical specialties where applications of supervised learning are more achievable and of lower risk—such as radiology or pathology—have already been impacted by AI. The US Food and Drug Administration has now approved over a dozen products solving medical tasks with AI [1].

Computer versus human intelligence

I agree with Eric Topol, when he stated: "Human health is too precious—relegating it to machines, except for routine matters with minimal risk, seems especially far-fetched" [1]. I would argue that concerns around AI taking over the jobs of physicians can be dispelled. Awareness, multi-tasking, flexibility and communication skills are human capabilities that no AI has achieved or seem likely to achieve anytime soon. Instead, I foresee that AI will remain in the co-pilot seat, improving our workflow and instilling more rationality into our practice. In particular, I am hopeful that significant developments are pending in the fields of real-time prediction of adverse events (patient deterioration, acute kidney injury, etc.), personalised drug dosing (antibiotics, etc.) and virtual scribes to help manage patient notes. Healthcare will not become fully automatized. A more realistic and achievable vision is represented by AI-driven decision support systems embedded in healthcare, capable of suggesting more optimal decisions (with confidence intervals) and reducing the daunting uncertainty we face in nearly every decision we make.

Electronic supplementary material
The online version of this article (https://doi.org/10.1007/s00134-019-05662-6) contains supplementary material, which is available to authorized users.

Compliance with ethical standards

Conflicts of interest
The author has no conflict of interest.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Received: 10 February 2019 Accepted: 7 May 2019
Published online: 24 June 2019

References

1. Topol EJ (2019) High-performance medicine: the convergence of human and artificial intelligence. Nat Med 25:44–56. https://doi.org/10.1038/s41591-018-0300-7
2. Bailly S, Meyfroidt G, Timsit J-F (2018) What's new in ICU in 2050: big data and machine learning. Intensive Care Med 44:1524–1527. https://doi.org/10.1007/s00134-017-5034-3
3. Ghassemi M, Celi LA, Stone DJ (2015) State of the art review: the data revolution in critical care. Crit Care 19:118. https://doi.org/10.1186/s13054-015-0801-4
4. Beam AL, Kohane IS (2016) Translating artificial intelligence into clinical care. JAMA 316:2368–2369. https://doi.org/10.1001/jama.2016.17217
5. Ince C (2017) Intensive care medicine in 2050: the ICU in vivo. Intensive Care Med 43:1700–1702. https://doi.org/10.1007/s00134-017-4808-y
6. Nemati S, Holder A, Razmi F et al (2018) An interpretable machine learning model for accurate prediction of sepsis in the ICU. Crit Care Med 46:547–553. https://doi.org/10.1097/CCM.0000000000002936
7. Pirracchio R, Petersen ML, Carone M et al (2015) Mortality prediction in intensive care units with the super ICU Learner Algorithm (SICULA): a population-based study. Lancet Respir Med 3:42–52. https://doi.org/10.1016/S2213-2600(14)70239-5
8. Héder M (2017) From NASA to EU: the evolution of the TRL scale in public sector innovation. Innov J 22:1–23
9. Gårdlund B, Dmitrieva NO, Pieper CF et al (2018) Six subphenotypes in septic shock: latent class analysis of the PROWESS Shock study. J Crit Care 47:70–79. https://doi.org/10.1016/j.jcrc.2018.06.012
10. Antcliffe DB, Burnham KL, Al-Beidh F et al (2018) Transcriptomic signatures in sepsis and a differential response to steroids: from the VANISH randomized trial. Am J Respir Crit Care Med 199:980–986. https://doi.org/10.1164/rccm.201807-1419OC
11. Seymour CW, Kennedy JN, Wang S et al (2019) Derivation, validation, and potential treatment implications of novel clinical phenotypes for sepsis. JAMA 321(20):2003–2017. https://doi.org/10.1001/jama.2019.5791
12. Sutton RS, Barto AG (2018) Reinforcement learning: an introduction, 2nd edn. MIT Press, Cambridge
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. 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
15. The House of Commons (2018) Algorithms in decision-making—Science and Technology Committee—House of Commons. https://publications.parliament.uk/pa/cm201719/cmselect/cmsctech/351/35102.htm. Accessed 11 Aug 2018