occupational specialist [such as an aviation medical examiner (AME)] and cardiologist.

Return to occupational duties may be considered based on a detailed risk assessment. Secondary prevention must remain optimal for all high-hazard employees.[2–5,13–16]

Conclusions

Cardiovascular disease is highly relevant in a population working in high-hazard occupational environments, whether civilian or military (such as aircrew, seafarers, submariners, special forces, bus drivers, or firefighters). The necessity for cardiac surgery will likely have a significant impact on both the lives and careers of these individuals. Therefore, it is critical that their assessments both pre- and post-surgery are detailed and specific as to ensure the best surgical outcome but also the minimal impact on their livelihoods.

Individuals who work in a high-hazard occupational environment may have pathology identified earlier than the so-called 'normal' population due to occupational cardiovascular screening. Furthermore, they often require careful long-term follow-up to ensure that they can continue to perform safely and that it is appropriate for them to do so, with a specific focus on the risks of distraction (that may result in loss of attention at a critical phase of work) or sudden incapacitation.

Whilst individuals should be treated using international guidelines, if more than one equivalent treatment approach exists, cardiac surgeons should consider which alternative is more appropriate given the occupational role undertaken; liaison with the specialist medical examiner (e.g. Occupational Health consultant or AME for aircrew) is strongly recommended prior to surgery to fully understand the extent of the impact of the cardiac surgery being considered on their future employment.

Many patients will be able to resume their activity post-operatively, albeit often with restrictions on their occupational duties. The post-operative follow-up of those at the highest occupational risk, such as aircrew, needs tight scheduling and requires seamless collaboration with the cardiologist and the occupational specialist (i.e. AME for aircrew).

Improving dialogue between the surgical and cardiological societies and the occupational and licensing authorities is self-evident to ensure support and collective agreement for future updates of the occupational health regulations and the clinical cardiac surgery guidelines.[17]

Conflict of interest: none declared.

References

References are available as supplementary material at *European Heart Journal* online.


Maj Thomas Syburra FRCS
Senior consultant cardiac surgeon
Luzern Kantonsspital - Herzzentrum
Spitalstrasse, 6000 Luzern 16, Switzerland.
Military Deputy Head Aviation Medicine
Swiss Air Force
Aeromedical Centre, Dübendorf, Switzerland
Tel: +41 41 205 47 14
Email: thomas.syburra@luks.ch


Norbert Guettler, MD, MA, FESC
Consultant Cardiologist & Electrophysiologist
Central Military Hospital, Koblenz, Germany
Colonel, Internal Medicine and Cardiology
Air Force Centre of Aerospace Medicine
Fürstenfeldbruck, Germany
Tel: +49 8141 5360 2033
Email: norbertguettler@bundeswehr.org


Joanna L. D'Arcy, MBChB MD FRCP DAvMed
Consultant Cardiologist
Dept. of Cardiology,
John Radcliffe Hospital,
Oxford, OX3 9DU, UK
Tel: +44 1865 220326
Email: Joanna.d'arcy@nhs.net


Edward D. Nicol, MD FRCP FRCR FESC FACC DAvMed
Corresponding Author
Consultant Cardiologist
Aviation Medicine Clinical Service
Centre of Aviation Medicine
RAF Henlow, Beds. SG16 6DN, UK
Tel: +44 1462 851515 ext. 6353
Email: e.nicol@nhs.net

doi:10.1093/eurheartj/ehz701

Artificial intelligence-assisted care in medicine: a revolution or yet another blunt weapon?

Potentials, challenges, and the future of implementing artificial intelligence (AI) for clinical care

Today, artificial intelligence is a sophisticated pattern recognition device

In the history of cardiovascular medicine, technical innovation has always been an essential driver for medical breakthroughs. An incomplete list could contain examples like Laennec's stethoscope, the electro- and echocardiogram, percutaneous coronary interventions, and continue to transcatheter structural heart interventions to open heart surgery, ventricular assist, and implantable electronic devices.

Nevertheless, we have also seen many examples of cutting edge technologies associated with big promises and bold expectations, that have not really translated to routine clinical care yet (e.g. cardiac cell therapy[1]). Recently, one can hardly overlook the work and exclamations regarding AI in medicine.

Artificial intelligence is neither a specific technology nor has it any 'artificial' features. More correctly referred to as machine intelligence

(MI), it describes a roughly 60-year-old research field of algorithmically solving problems that conventionally require human intelligence.[2] The MI field has seen times when expectations were disappointed when the field was even announced to be dead—such episodes emerged multiple times in the history of MI and some periods were even called the winters of MI, accompanied by periods of dramatically reduced interest and funding. Multiple factors caused them, but the most important cause was over-inflated expectations that were doomed to be disappointed.

We are currently witnessing the third advent of MI, mainly driven by the application of deep neural nets to large amounts of data. Just recently, remarkable advances were achieved in image and language processing, reaching human parity in specific applications, such as image classification[3] or machine translation.[4,5] However, one has to keep in mind that MI shines in specific use cases [i.e. spot cancer on the computed tomography (CT) scan], but that is not usually what one has in mind when thinking about general AI (i.e. 'Hollywood's robots'). Most people associate the term AI with a machine's ability to mimic human cognitive function. A vision that factually is still far out of reach.

The most effective existing MI techniques augment and support our own intelligence. Most recent implementations of MI are based on deep neural network architectures,[6] which are basically complex statistical pattern recognition machines. Nevertheless, deep learning is a powerful tool with great potential for use in medical practice and research. Among the more common deep learning applications for diagnostic support (i.e. improving the quality of specific clinical tasks), some research has also led to new findings. For instance, Poplin *et al.*[7] were able to infer patients' gender as well as cardiovascular risk factors only using fundoscopic images. Attia *et al.*[8] used deep learning to detect the electrocardiographic signature of atrial fibrillation present during normal sinus rhythm using just a standard 10 s, 12-lead electrocardiogram. We ourselves at the German Heart Center, Berlin, developed highly accurate deep learning-based models for detecting complications during post-cardiothoracic surgery care.[9] A recent review presents many more examples of impressive applications of deep learning in the cardiovascular field.[10]

Clinical translation of research results—the final frontier

However, despite such exciting academic deep learning achievements, Panth *et al.*[11] state that 'at present, the algorithms that feature prominently in the research literature are in fact not, for the most part, executable at the frontlines of clinical practice'. Translating MI research results into clinical care is a daunting task, for which researchers do not get much credit. However, since mid of 2018, we are in the process of translating our results into clinical care. We developed a prototype medical product and integrated it into the hospital's medical IT infrastructure for research use (*Figure 1*).

One experiences countless obstacles on the way to a certified medical product, specifically in the context of hospitals' individual data infrastructures and regulations. Hospitals are still undergoing a transition to holistically digitized units; thus, access to digitized data is limited. Furthermore, historically, clinical information systems have been organized into multiple orthogonal data silos, making data harmonization considerably more challenging. Also, hospitals adopt data structures to their individual requirements, and rarely make use of coding and terminology standards such as LOINC[12] or Snomed CT.[13] Adapting a hospital's established infrastructure to these standards in order to integrate a new device or system can be time-consuming and costly.

Figure 1 Photo of the current prototype of the post-operative bleeding prediction algorithm based on the results of Ref.[9] Here, all patients of the ward are displayed with corresponding risk values calculated in real-time. The current value and the course of the past 3 h are displayed graphically for each patient; patients at risk are highlighted.

Since the introduction of the new EU Medical Device Regulation (MDR) from 2017, becoming effective in 2020, broader ranges of software types are considered as medical products and, additionally, all instances of software as a medical device (SaMD) received higher risk classes, coupled with more complicated, time-consuming, and costly requirements for certification. Such raised risk classes also apply to clinical decision support systems, which intend to *monitor physiological processes* or *provide information which is used to take decisions with diagnosis or therapeutic purposes.*[14]

The EU's General Data Protection Regulation (GDPR[15]), effective since 2018, has introduced additional challenges related to health data access. While clinical research may use pseudonymized data, the development of medical products requires the use of either fully anonymised data (which is difficult, if not impossible, for datasets with a large variety of values) or project-specific consents from the data owner (i.e. patients).

Still, it is worth to tackle all these challenges and translate deep learning achievements into clinical use.

Deep learning—the only kid on the block?

There are two fundamental types of models:

(1) entirely data-driven models learned from empirical observations, and
(2) knowledge-based models which build upon prior insights.

Deep learning belongs to the class of non-knowledge-based approaches and is favourable when large quantities of data are available. On the other hand, when

(1) data for learning is difficult to acquire,
(2) the system needs to follow specific rules or respect boundaries (e.g., guidelines and risk-benefit ratio), or
(3) clinicians need complete computational transparency and reproducibility (e.g., to study a system's behaviour),

then knowledge-based approaches are more suitable. Various knowledge-based methodologies exist differing in the complexity of modelling and computation; from simple decision trees, through rule-based approaches (e.g. Arden Syntax[16]), to more complex probabilistic graphical models (e.g. Markov random fields[17] and Bayesian networks[18,19]).

These methodologies require a predefined knowledge base, which can be learned automatically from structured data, semi-automatic, or completely modelled by domain experts. For instance, when only a small amount of data is available, domain experts can support the modelling by defining the usually known causal relationship between variables, and the model learns only the probabilities.

While the main effort in non-knowledge-based approaches lies mainly in data preparation (which may decrease in the future through appropriate data collection), modelling a knowledge base is usually very time intensive and domain experts require a deep understanding of the underlying methodology. Nevertheless, once a knowledge base exists, it is usually human-readable and, therefore, allows experts to review it and modify individual decision-making characteristics (i.e. for updating, 'fine-tuning', or clinical adjustments). For instance, a probabilistic graphical model can express a comprehensive view of clinical guidelines and related statistics[20] and, in the next step, run inferencing algorithms to compute model-based medical evidence.[21]

Building sophisticated knowledge bases is a process in which medicine has gained considerable experience: developing resources such as the ESC/EACTS joint guideline for valvular heart disease[22] is a prime example. Therefore, we suggest that creating such knowledge bases should become a substantial part of defining clinical guidelines or even earlier when conducting clinical trials.

Another important class of MI models using prior knowledge incorporate biophysical properties. For instance, computational fluid and airways simulations[23] are applied on radiological image sequences in order to visualize and simulate dynamic processes for diagnostic as well as potential effects of treatment options.[24,25]

Many opportunities for computer-support in healthcare appear, and individual MI methodologies have unique advantages and

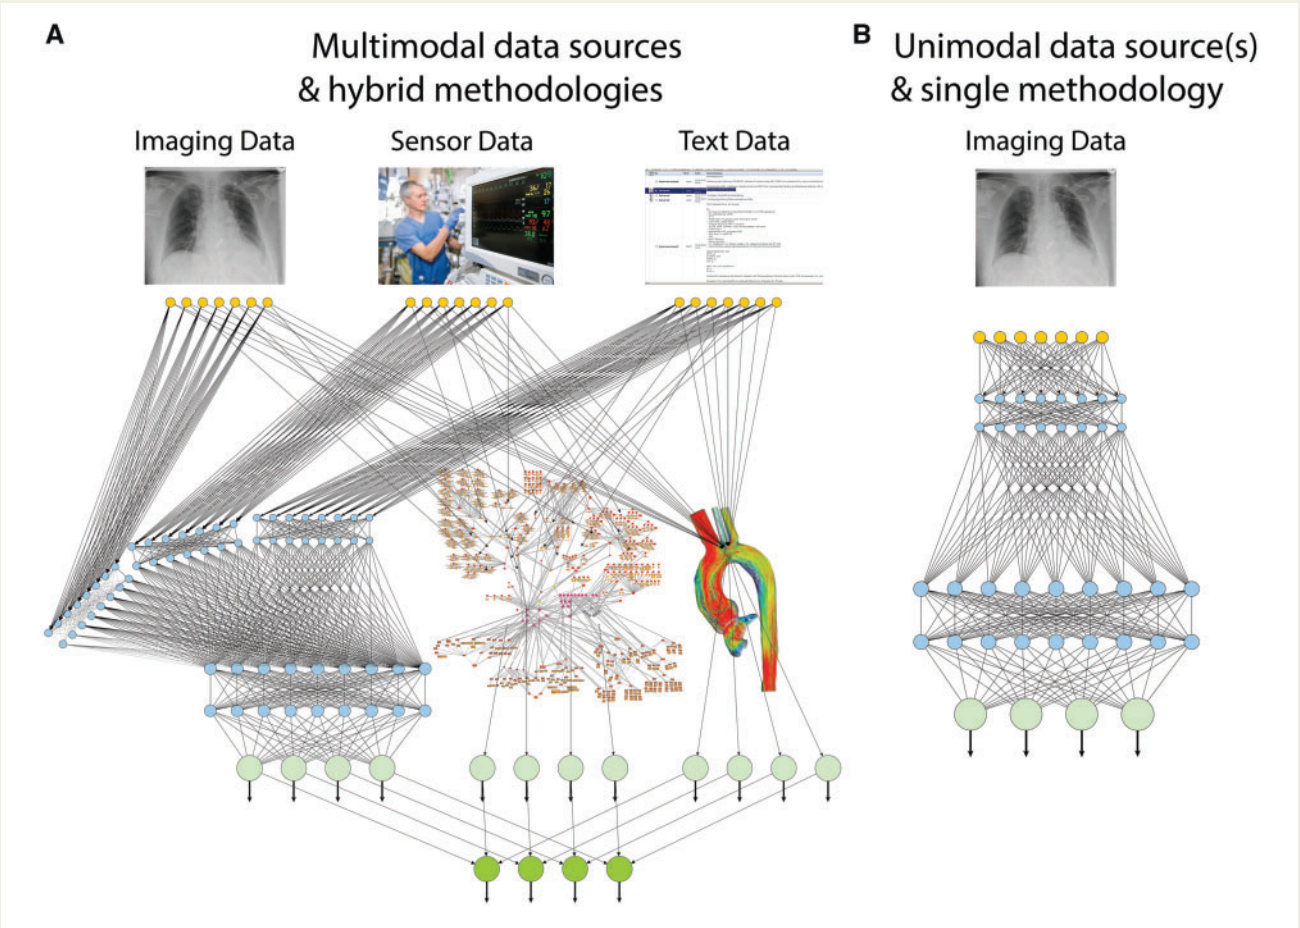

Figure 2 (*A*) An exemplary illustration of different input modalities (images, time series, and diverse structured data, such as genomic sequence data[26]) and modelling methodologies, such as deep neural networks, probabilistic graphical models, computational fluid dynamics. These approaches can be combined into 'one model' referring to hybrid modelling or server for inter-model validation. In contrast, (*B*) shows a single modelling methodology, in this case, a deep neural network, with a unimodal data input.

disadvantages; therefore, most have good chances to be further developed and clinically applied. Combined methodologies are promising to increase the range of clinical applicability as well as increase their accuracy (see *Figure* 2*A* and *B*).

A likely future of machine intelligence in medicine

Machine intelligence, mainly in the form of deep learning, is here to stay! Its statistical learning nature, however, may pose problems for medical application,[27] but these can be overcome by making use of explainable functionality on purely statistical models, and by combining the power of knowledge-based approaches with deep learning. Besides, MI performance on medical data will likely see a significant increase in accuracy and robustness when using multi-modal data as inputs.[26] *Figure* 2*A* illustrates this by showing a deep neural network with different input data modalities. Within the next decade, we will see a fundamental shift in clinical research and practice. It is up to physicians to shape this transition in order to avoid a repetition of the electronic health record introduction, which has happened mainly without physicians' engagement and may have resulted in a plain usability catastrophe for this very reason.[28]

Generally speaking, MI development has the potential to disrupt healthcare systems and clinical care: computers can process large quantities of data and structured representation of knowledge in just a short time without loss of information. They are tireless, and without personal interests. The digitization and desire for personalized medicine are likely to establish new clinical domains focusing on computer-assisted medicine. Indeed, machines will take over, especially for recurring and specific foreseeable tasks; however, this is a unique opportunity to regain more time for patient communication, sensitive matters, and complex decision-making.

Machine intelligence holds great promise to the third world by enabling ubiquitous access to expert-level diagnosis. In a digitized and machine-supported healthcare setting, it is conceivable that machines may receive a situation-specific autonomy, which will depend on the combination of both (i) a patient's mortal danger and (ii) health professionals' required time to intervene.

Situation-specific autonomy is particularly reasonable in extreme and isolated situations, e.g. in catastrophes on Earth as well as during manned deep-space exploration.

Finally, since MI-based Research & Development is increasing exponentially fast, stronger regulations for data access, and medical devices need fundamental operational rethinking and restructuring. For instance, international initiatives break the mode of data sharing to make data FAIR[29]: Findable, Accessible, Interoperable, and Reusable.

Acknowledgements

A.M. is participant in the BIH Charité Clinician Scientist Program funded by the Charité – Universitätsmedizin Berlin and the Berlin Institute of Health.

Funding

C.E. is funded by the Swiss National Science Foundation (SNSF) Ambizione Program (grant agreement number 174025).

†The first two authors contributed equally to the study.

Conflict of interest: A.M. declares the receipt of consulting and lecturing fees from Medtronic and consulting fees from Pfizer. C.E. is a co-founder and holds shares of codiag AG. All other authors have no conflict of interest to declare.

References

References are available as supplementary material at *European Heart Journal* online.

Corresponding Author:
Alexander Meyer†, MD
Department of Cardiothoracic and Vascular Surgery
German Heart Center Berlin
Augustenburger Platz 1,
13353 Berlin, Germany
Email: meyera@dhzb.de
Tel: +49 (0)30 4593 1000
Fax: +49(0)30 4593 2100

Mario A. Cypko†, PhD
Human Research Office,
Directorate of Human Spaceflight and Robotic Exploration
Programmes, European Space Agency,
Keplerlaan 1, 2200 AG Noordwijk,
The Netherlands

Carsten Eickhoff, PhD
Center for Biomedical Informatics,
Brown University, 233 Richmond Street,
Providence, RI 02912, USA

Volkmar Falk, MD
Department of Cardiothoracic and Vascular Surgery, Deutsches Herzzentrum Berlin, Augustenburger Platz 1, 13353 Berlin, Germany;
Department of Cardiothoracic Surgery, Charité – Universitätsmedizin Berlin, Charitéplatz 1, 10117 Berlin, Germany

Maximilian Y. Emmert, MD, PhD
Department of Cardiothoracic and Vascular Surgery, Deutsches Herzzentrum Berlin, Augustenburger Platz 1, 13353 Berlin, Germany;
Department of Cardiothoracic Surgery, Charité – Universitätsmedizin Berlin, Charitéplatz 1, 10117 Berlin, Germany