

Artificial Intelligence for the Artificial Kidney: Pointers to the Future of a Personalized Hemodialysis Therapy

Miguel Hueso^a Alfredo Vellido^b Nuria Montero^a Carlo Barbieri^c
Rosa Ramos^c Manuel Angoso^d Josep Maria Cruzado^a Anders Jonsson^e

^aDepartment of Nephrology, Hospital Universitari Bellvitge, and Bellvitge Research Institute (IDIBELL), L'Hospitalet de Llobregat, Spain; ^bIntelligent Data Science and Artificial Intelligence (IDEAI) Research Center, Universitat Politècnica de Catalunya (UPC), Barcelona, Spain; ^cFresenius Medical Care, Bad Homburg, Germany; ^dDialysis Unit, Clínica Virgen del Consuelo, Valencia, Spain; ^eArtificial Intelligence and Machine Learning Research Group, Universitat Pompeu Fabra (UPF), Barcelona, Spain

Keywords

Artificial intelligence · Artificial kidney · Hemodialysis · Machine learning · Patient safety

Abstract

Background: Current dialysis devices are not able to react when unexpected changes occur during dialysis treatment or to learn about experience for therapy personalization. Furthermore, great efforts are dedicated to develop miniaturized artificial kidneys to achieve a continuous and personalized dialysis therapy, in order to improve the patient's quality of life. These innovative dialysis devices will require a real-time monitoring of equipment alarms, dialysis parameters, and patient-related data to ensure patient safety and to allow instantaneous changes of the dialysis prescription for the assessment of their adequacy. The analysis and evaluation of the resulting large-scale data sets enters the realm of "big data" and will require real-time predictive models. These may come from the fields of machine learning and computational intelligence, both included in artificial intelligence, a branch of engineering involved with the creation

of devices that simulate intelligent behavior. The incorporation of artificial intelligence should provide a fully new approach to data analysis, enabling future advances in personalized dialysis therapies. With the purpose to learn about the present and potential future impact on medicine from experts in artificial intelligence and machine learning, a scientific meeting was organized in the Hospital Universitari Bellvitge (L'Hospitalet, Barcelona). As an outcome of that meeting, the aim of this review is to investigate artificial intelligence experiences on dialysis, with a focus on potential barriers, challenges, and prospects for future applications of these technologies. **Summary and Key Messages:** Artificial intelligence research on dialysis is still in an early stage, and the main challenge relies on interpretability and/or comprehensibility of data models when applied to decision making. Artificial neural networks and medical decision support systems have been used to make predictions about anemia, total body water, or intradialysis hypotension and are promis-

Miguel Hueso, MD
Department of Nephrology
Hospital Universitari Bellvitge and Bellvitge Research Institute (IDIBELL)
C/ Feixa llarga, s/n, ES-08907 L'Hospitalet de Llobregat, Barcelona (Spain)
E-Mail mhueso@idibell.cat

Alfredo Vellido, PhD
Intelligent Data Science and Artificial Intelligence (IDEAI) Research Center
Universitat Politècnica de Catalunya (UPC)
C/ Jordi Girona, 1–3 ES-08034 Barcelona (Spain)
E-Mail avellido@cs.upc.edu

ing approaches for the prescription and monitoring of hemodialysis therapy. Current dialysis machines are continuously improving due to innovative technological developments, but patient safety is still a key challenge. Real-time monitoring systems, coupled with automatic instantaneous biofeedback, will allow changing dialysis prescriptions continuously. The integration of vital sign monitoring with dialysis parameters will produce large data sets that will require the use of data analysis techniques, possibly from the area of machine learning, in order to make better decisions and increase the safety of patients.

© 2018 S. Karger AG, Basel

Background

Renal transplantation is the best option for end-stage renal diseases in terms of patient life expectancy and quality of life, but the shortage of organ donors and clinical contraindications make dialysis the only real option for many patients [1]. Complications in dialysis occur during the treatment, but current dialysis software is not able to make real-time changes when unexpected events happen or to learn from experience to improve therapy and personalize treatments. Great efforts have been made to preset ultrafiltration (UF) and dialysate sodium profiles to counterbalance the negative effects of uncontrolled water or solute removal [2]. However, this approach has failed since predetermined UF and sodium profiles do not allow adaptations during treatment if the designed profile was inadequate [3]. Thus, we need real-time data monitoring using multiple biosensors to detect the most important chemical/physical signals from patients and achieve an automatic, immediate and adequate change of dialysis parameters to reduce patient morbidity [4].

Data analytics, delivered by artificial intelligence (AI) and machine learning (ML), is likely to play an important role in patient monitoring of dialysis efficacy and safety. AI is already having an impact on health care in areas such as medical image analysis, smart robotics in surgery, and voice-enabled assistants [5, 6]. Can this success be exported to dialysis? Is it possible to design and develop smart dialysis devices? A first scientific meeting to discuss the present and future of AI in dialysis was organized at the Hospital Universitari Bellvitge (L'Hospitalet, Barcelona) with the aim to answer these questions and discuss the state of the art in the field. As a result from that meeting, this study reviews AI experiences in dialysis and discusses barriers, challenges, and future applications in the field.

Present Artificial Intelligence in Medicine: Challenges for the Application of Machine Learning in this Area and the Importance of Model Interpretability

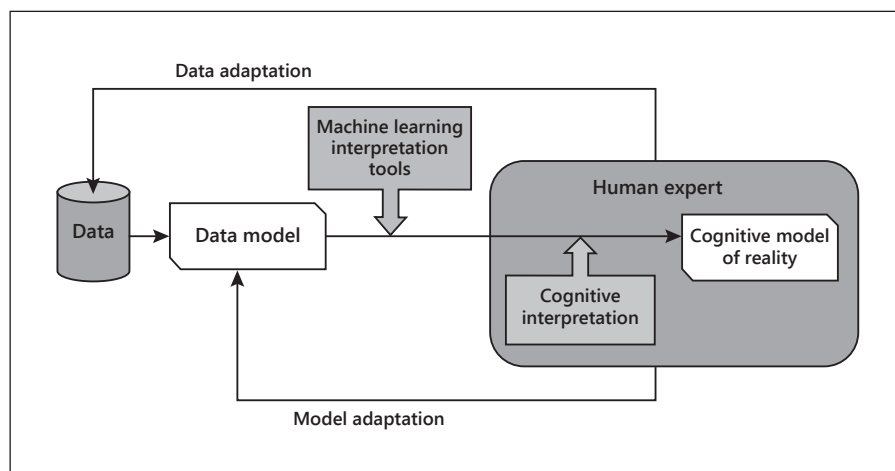
Current research in the life sciences relies heavily on data acquisition and analysis [7]. The drive towards data-based science is especially strong in the -omics fields, where fast technological advances in data acquisition have shifted part of the research challenges to the computer science domain. These challenges may be only relatively less pressing in health care or biomedicine, but, even here, the complexity and heterogeneity of medical information mean that “it is not yet possible to create a comprehensive model capable of considering all the aspects of health care systems” [8].

The widespread adoption of computers and networked digital systems is creating a fast-evolving medical information ecosystem. This situation of data abundance, only bound to increase, may be seen as the perfect opportunity for data analytics, be it in the form of statistics, ML, or combinations of both under the umbrella of AI. Data analytics is at the heart of medical decision support systems (MDSS), which are often based on AI methods such as ML. MDSS, although still far from the mainstream of medical practice, have made substantial advances in specific domains [9–11].

This opportunity seems like a win-win scenario but, in fact, the implementation of AI methods in health care and biomedicine is riddled with nontrivial challenges. We argue here that, unless those challenges are appropriately addressed, this type of methods is unlikely to be adopted in practice beyond a limited number of niche applications. Some of these pressing challenges have their origin in societal issues. Let us illustrate this with 2 examples. One of the currently most fruitful applications of ML in medicine is electronic health records (EHR) text mining. It has been argued, though, that this application might lead to a reduction of skills among medical experts. This negative consequence of the use of ML methods in medicine has been described as ML methods' undue “focus on text and the demise of context” [12]. The second example involves the implementation of a European Union directive for general data protection regulation that will be enforced in 2018 and mandates a *right to explanation* of all decisions made by automated or AI algorithmic systems [13]. Note that this regulation should be a warning for any AI-based MDSS for which the required explanation is unfeasible.

Let us now turn to a challenge that is directly derived from the characteristics of many AI methods: the poten-

Fig. 1. Interpretability cycle, considering a number of *actors*, including the data and/or corresponding models and the ML interpretation tools, on one side (the *machine* side), and the cognitive processing based on a model of the reality to be interpreted, on the other side (the *human* side). The cycle allows for data and model adaptation according to human interpretation (adapted from Vellido et al. [14]).



tial lack of interpretability and/or comprehensibility of the data models they generate. Interpretability has become a central issue in ML research over the last few years [14]. Such a surge of interest is at least partly caused by the resurgence of connectionist ML in the form of deep learning, a family of successful methods that have also found their way into biomedicine and health care [15].

Three main challenges for the application of ML in medicine have recently been listed [12], and one of them is precisely interpretability, expressed as “the need to open the machine learning black box.” Not that this is a new challenge for ML, because the *black box syndrome* was already on the table decades ago [16]. The problem can also be expressed as the impossibility to describe in clear terms the relationship between the observed data and the resulting outcomes due to its complexity. This has obvious implications in the medical context: if an MDSS churns out decisions that cannot easily be described in comprehensible terms, a potentially insurmountable barrier is raised between the MDSS and the human users. For instance, the medical expert could not trust to implement a decision that cannot be explained, while the patient might not trust an expert that bases her or his judgment on unexplainable algorithmic outcomes. Also, and on the basis of legal safeguards such as the general data protection regulation, a health care system might not be willing to implement an opaque MDSS in clinical practice.

It has been argued that the role of ML in health care should be acting “as a tool to aid and refine specific tasks performed by human professionals” [17]. Note that this means that interpretability cannot be dissociated from the human interpreter. Some formal framework for ma-

chine-human interaction in the pursuit of interpretability is required. One such framework is outlined in Figure 1.

A similar augmented framework has recently been proposed for ML interpretability through visualization [18], emphasizing interactivity between human and algorithm. Visualization can be seen as a powerful tool for exploratory data analysis, and it has also been mentioned to play a central role in interpretability for medicine in recent research [19].

This is the context in which we need to appraise the present application of AI and ML in nephrology and, specifically, dialysis. Research in this area is still somehow disconnected and tentative. For instance, AI and ML have addressed problems concerning anemia and other issues in hemodialysis (HD) patients [20, 21]. Artificial neural networks were used for predicting total body water in HD patients [22] and the target range of plasma intact parathyroid hormone concentration [23]. Support vector machines and reinforcement learning have been proposed for erythropoietin dosage prediction and personalization [24]. Full ML-based MDSS have also been proposed, like a tool for anemia management in renal disease patients undergoing dialysis [25].

In continuous ambulatory peritoneal dialysis for monitoring patients with severe kidney failure, ML has been used as the basis for an MDSS for blood test analysis to ascertain their stroke risk levels [26]. A completely different use of ML is the application of classification trees and naïve Bayes to the prediction of the quality of life in HD patients [27]. Data feature selection can help the interpretation of a given problem by providing guidance about the relative impact of input features on the outcome. An example of such an approach is the analysis of the relative

relevance of genetic and phenotypic data features associated with the inflammatory status of patients on dialysis [28].

As with the application of AI and ML techniques to medicine in general, a word of caution about the existence of nontrivial challenges must be included here also for the case of dialysis. Ricci et al. [29] have recently described a scenario where “often the dialysis/continuous renal replacement machine is seen as a black box where nothing can be modified from the scheduled functionality and the machine is automatically selecting the best operational mode to achieve therapeutic targets.” Authors warn that this scenario might lead to the progressive deskilling of clinicians. It is also a context in which lack of model transparency can only be overcome by achieving a synergy between machine and human experts in which both tap into each other’s expertise to increase their own.

Clinical Experiences of AI and ML Techniques to HD: Anemia Therapy Control and Hemocontrol™

Anemia Therapy Control, or How Algorithms Make Systems Smart

Although dialysis partially restores kidney blood filtration function, it is unable to replace its role in regulating metabolism and in endocrine function, producing common dialysis complications such as anemia [30]. The availability of exogenous erythropoiesis-stimulating agents (ESA) has improved the treatment of anemia but a high intra- and interindividual response variability has been detected [31]. Thus, with the aim to support the nephrologist in ESA and iron dosing, an AI tool has been developed. The anemia control model (ACM) consists of 2 components: (1) an artificial neural network model that uses updated patient clinical data to predict future hemoglobin concentration; (2) an algorithm that suggests the optimal ESA and iron dosage to achieve the hemoglobin target. It is important to stress that an ACM only provides therapy recommendations and that has to be validated.

The neural network was trained to predict the change in hemoglobin value occurring over a month, based on the current patient condition (lab tests, baseline characteristics, and dialysis treatment parameters) and the administered ESA and iron doses. This predictive model is used to assess how hemoglobin is expected to change as a function of different ESA dosages, and is used in the next dose selection step [32, 33]. In summary, based on model simulation, the algorithm looks for the drug dosage that would move hemoglobin to the target interval, while

avoiding excessive hemoglobin drops or jumps [34]. Patients’ updated clinical information is fed to the ACM by means of an automatic interface module with the clinical system, which ensures continuous recording of clinical and laboratory data.

The clinical introduction of the ACM has allowed an increase in the number of patients in target (70.6–83.2%), a decrease in hemoglobin variability (from 0.95 to 0.83 g/dL), a significant reduction of hemoglobin >12 g/dL (18.06–7.5%) and a reduction of ESA and iron consumption [25].

Hemocontrol™ and ML Systems

Intradialytic hypotension (IDH) occurs in 20% of patients and increases cardiovascular mortality and morbidity [35]. IDH are caused by an imbalance between the prescribed weight loss, the plasma refilling capacity, and the cardiovascular system compensatory mechanism. To prevent IDH, an innovative, multi-input multi-output controller named Hemocontrol™ has been developed [36]. Hemocontrol™ assesses 3 variables (blood volume, total weight loss, and average dialysate sodium level) and controls 2 dialysis parameters (UF and dialysate sodium level). Relative blood volume (RBV) changes are monitored by an optical absorption biosensor (Hemoscan™) that estimates hemoglobin concentration using spectrophotometry. The estimated RBV used on the hemocontrol prescription is based on the patient’s ratio of total UF and final blood volume changes (blood volume/UF volume). Hemocontrol™ modifies sodium concentration of dialysate and regulates the UF rate through a biofeedback mechanism to adjust it to the predetermined RBV trajectory. The changes in UF rate and in sodium concentration are carried out within a defined limit to avoid sodium overload and to achieve the prescribed weight loss. Biofeedback systems that only adjust UF without changing dialysate sodium concentration will not reduce the rate of IDH compared to conventional therapy [37]. The dialysis team will be able to follow up the patient’s response adjusting the UF goal and change of RBV as required by the patient’s clinical condition.

Biofeedback systems like Hemocontrol™ have been shown to significantly reduce the incidence of IDH by 39% [38] and myocardial stunning [39] while achieving higher UF volumes per treatment [40] compared to conventional HD. The Hemocontrol™ initial rise in the dialysate sodium concentration has been shown to increase the plasma levels of vasopressin that might contribute to

a decrease in the number of IDH episodes at the end of the treatment [41]. It has also been shown to reduce the number of antihypertensive medications in the long term, improving the patient's HD experience by decreasing post-HD fatigue, increasing the number of event-free sessions and reducing the number of dialysis sessions that require nursing interventions [42].

The Future in Medical Devices That Will Revolutionize the Treatment of Dialysis Patients

Miniaturization, portability, flexibility, "water use" efficiency, and wearable technology are goals subject to intense research that will contribute to the metamorphosis of current dialysis machines. Achieving these goals will require innovative technological developments in the field of: membranes (smart biocompatible nanotechnology-produced membranes or composite membranes containing specific sorbents) [43]; dialysis fluid regeneration with cation exchange sorbent systems and enzyme technology [44]; highly efficient pumping systems both for blood and dialysis fluids; anticoagulation and non-thrombogenic surfaces for clotting avoidance, and a safe vascular access for a therapy that should operate continuously [45]. Developments in bioelectronics and semiconductor technologies have made possible to design circuits that can be integrated into a single chip and consume less power. Smaller but powerful batteries with longer backup time will further allow miniaturization, enhancing portability. This technology has allowed the development of the wearable artificial kidney based on dialysate-regenerating sorbent technology [46]. Another development in miniature dialysis devices, still in the preclinical study stage, is the implantable bioartificial kidney based on silicone nanotechnology to create membranes mimicking glomerular filtration in combination with renal tubular cells [47].

Wireless technology will also enhance clinical workflow for EHR or remote monitoring. However, electrical leakage in wearable devices is a matter of concern given that patient safety is a key challenge for the development of future dialysis devices, where stringent regulatory and compliance requirements catalyze innovation. To improve patient security, a system of alarms is required for air detection, blood leaks, transmembrane pressure, and proper blood pump functioning. In addition, miniaturized sensor technology in dialysate affluent and effluent are also important for adequate monitoring of electrolytes and acid-base disturbances. Online monitoring sys-

tems with automatic instantaneous biofeedback based on real-time and repeated measurement of chemical and physical signals from the patient, such as blood volume changes, dialysate conductivity, urea kinetics and thermal energy balance, followed by the analysis and evaluation of the data, will allow more personalized treatments [4]. Treatment parameters will change continuously through a dialysis delivery system that incorporates adaptive and logic controls. Integrating monitoring vital signs with parameters of dialysis will produce large data sets that will require predictive models with complex calculations and, possibly, ML methods to deliver better decisions and increase the patients' safety [48, 49].

Is AI the Future for Therapy and Decision Problem Solving in Medicine or Is It Science Fiction? A Vision of Potential Applications of AI in Medicine

To discuss the future potential of AI in medicine it is necessary to assess its current standing, as described in the previous sections. AI is the field of computer science that studies the automatic generation of intelligent behavior from a computational point of view. The term "intelligent behavior" is usually defined in terms of how difficult it would be for a human to perform a given task [50]. Computational problems that are historically considered part of AI include reasoning, knowledge discovery, planning, learning, natural language processing, perception, and the ability to move and manipulate objects.

ML, a subfield of AI, has experienced an unprecedented boom in the 2010s, particularly in the form of deep learning [51]. This progress is intimately tied to the increased computational power of processing units. Many of the algorithms popularized during this period were developed decades earlier, but their full potential was not realized until the advent of commercially available graphical processing units that made it possible to efficiently parallelize computation. Automatically generating intelligent behavior that is comparable to that of humans has been a long-standing goal of AI. There are many recent publications that claim to achieve or surpass human-level performance in various areas of AI, such as image segmentation [52], concept learning [53], and speech recognition [54]. With respect to medicine, achieving human-level performance is not only of academic interest, but effectively determines when an AI algorithm may perform its task better than a human would. Several studies indicate that AI algorithms may already surpass humans in clinical tasks such as predicting lung cancer [55], skin

cancer [56], and the risk of heart attacks [57]. This trend is bound to continue in the future, with some clinical tasks performed more efficiently and accurately using AI.

With this background in mind, it is evident that AI holds a great potential for clinical tasks in medicine. However, much research on AI for medicine has been carried out in academic settings. For AI systems to have an impact on the diagnosis and treatment of the general population, it would be necessary to implement and integrate such systems in the public health care domain. As we discuss below, such a large-scale implementation of AI systems in medicine is far from straightforward. The problem of interpretability has already been discussed, but there is also a host of other issues that we describe below.

A prerequisite for the implementation of AI systems in medicine is to ensure that the appropriate infrastructure is in place. Specifically, patient data should be properly collected, labeled and organized, and measurements performed by digital tools that are directly connected to the network, such that patient records are automatically updated and maintained. Such a digital reform is under way in many countries [58], but there is still plenty of work to be done. An important limitation of ML algorithms is that they need access to large amounts of high-quality data. Although open access to data is becoming more common in academics, there are many incentives in medicine for keeping data private: companies want to maintain an edge over their competition; hospitals may not want to compromise the privacy of their patients, etc. Most likely, it would be necessary to propose and implement new regulations regarding how medical data are created, maintained and shared. A related problem is to ensure that sensitive patient information is not divulged in data records used to train ML algorithms. Patient records are usually anonymized to leave out sensitive information in order to guarantee the privacy of patients. However, as ML algorithms become increasingly powerful, they can also be used to predict the missing information in patient records. Hence, anonymization may not be sufficient on its own, and additional techniques are necessary to ensure that anonymized patient information cannot be reconstructed.

If AI algorithms are used to make clinical decisions, this raises both ethical and legal questions. Who is liable in case something goes wrong, the developer of the AI system or the hospital? What should an AI algorithm do in case it is faced with 2 unattractive options? Ethical dilemmas of this type are frequently discussed in the context of self-driving cars [59], e.g. what should an auton-

ous car do if it is faced with the decision of hitting a pedestrian or hurting its own passengers? The Ethics Commission at the German Ministry of Transport and Digital Infrastructure has recently published the world's first ethical guidelines for autonomous driving [60], and similar initiatives would be needed for medicine.

If discriminatory bias is present in historical data records, then ML algorithms simply learn to reproduce this discriminatory bias. There are several examples in the news of ML algorithms that discriminate: an algorithm implemented by Google to display job ads presented men with higher-paying job offers than women [61], and an algorithm for predicting crime rates in the USA unfairly associated a higher risk with African-Americans than with other ethnical groups [62]. It is easy to imagine clinical situations in which decisions may be considered discriminatory, so this is another aspect of AI that has to be taken into account.

Another problem is that implementing AI systems in health care can be viewed by doctors as *trespass* on their domain of expertise, and perhaps even as a threat to their job security. Given the current state of the art in AI, most doctors should be safe: AI algorithms excel at specialized tasks for which they have been trained, but they lack the ability to draw conclusions, integrate information and perform high-level reasoning. In addition, doctors can provide patients with a sense of comfort and human empathy that would be very difficult to reproduce in an artificial system. However, there are specialist doctors such as radiologists whose job mainly consists in analyzing images to diagnose patients. These specialists would directly compete with AI algorithms that perform the same task. Even in these cases, however, there is a danger in relying too much on AI algorithms, since the human knowledge might eventually become lost, an effect known as "automation bias."

In the short term, we believe that the proper way to integrate AI in medicine is to view AI algorithms as tools that doctors can use to make more informed decisions. In this sense, it is important to develop applications that present the outcome of AI algorithms in an explanatory way that doctors can understand and interpret. On the one hand, doctors would have to accept that there are certain tasks that AI algorithms perform better. On the other hand, doctors would be ultimately responsible for the final decisions taken regarding diagnosis and treatment.

It is difficult to speculate in how AI will transform medicine in the long term. One area of AI called life-long learning studies the problem of AI systems that evolve over time and adapt to new challenges and tasks that they

were not originally designed for. Such a system would more closely resemble a human's ability to switch between tasks and draw on previous experience to solve new tasks more efficiently. This research is still in an early stage, and it will likely be many years before such systems are ready to be implemented in real-world applications. One can imagine deploying a team of autonomous nanobots in a human body that automatically learn to perform new tasks over time, with the aim of ensuring the well-being of the patient.

Conclusions and Future Trends

End-stage renal disease patients depend on technology for living. The evolution of dialysis therapy has been characterized by the search for safer devices with more efficient and clinically tolerable treatments.

Data analytics, delivered by AI and ML, is likely to play an important role in medical decision making for patient monitoring of dialysis efficacy and safety. However, MDSS are still far from the mainstream of medical practice.

One challenge for AI methods is the potential lack of interpretability and/or comprehensibility of the data models they generate due to the impossibility to describe the relationship between the observed data and the result-

ing outcomes due to its complexity. This problem can only be overcome by achieving a synergy between machine and human experts in which both tap into each other's expertise to increase their own.

Acknowledgment

We thank Estanislao Navarro for the critical reading of the manuscript.

Conflict of Interest Statement

The authors have no conflicts of interest to disclose.

Author Contributions

This report comes from the 1st Science for Dialysis Meeting: Artificial Intelligence, which took place in the Hospital Universitari Bellvitge on September 22, 2017. M.H. and N.M. wrote the section on the background and future in medical devices. As chairmen they oversaw the consistency of the manuscript edition. Other sections were written by the following participants: A.V. wrote on the present AI in medicine, C.B. and R.R. on anemia therapy control, M.A. on Hemocontrol™, J.M.C. on the future in medical devices, and A.J. wrote on AI as the future for therapy and decision problem solving in medicine. Conclusions were written by all participants.

References

- 1 Saran R, Robinson B, Abbott KC, Agodoa LYC, Ayanian J, Bragg-Gresham J, et al: US renal data system 2016 annual data report: epidemiology of kidney disease in the United States. *Am J Kidney Dis* 2017;69:A7–A8.
- 2 Fichoux A, Gayraud N, Duranton F, Guzman C, Szwarc I, Vetromile F, Brunet P, Servel MF, Argilés A: A reliable method to assess the water permeability of a dialysis system: the global ultrafiltration coefficient. *Nephrol Dial Transplant* 2017;32:364–370.
- 3 Ronco C, Ghezzi PM, La Greca G: The role of technology in hemodialysis. *J Nephrol* 1999; 12(suppl 2):S68–S81.
- 4 Locatelli F, Buoncristiani U, Canaud B, Köhler H, Petitclerc T, Zuchelli P: Hemodialysis with on-line monitoring equipment: tools or toys? *Nephrol Dial Transplant* 2005; 20:22–33.
- 5 Senders JT, Arnaout O, Karhade AV, Desenbrock HH, Gormley WB, Broekman ML, Smith TR: Natural and artificial intelligence in neurosurgery: a systematic review. *Neurosurgery* 2017, Epub ahead of print.
- 6 Feng R, Badgeley M, Mocco J, Oermann EK: Deep learning guided stroke management: a review of clinical applications. *J Neurointerv Surg* 2017, Epub ahead of print.
- 7 Leonelli S: *Data-Centric Biology: A Philosophical Study*. Chicago, University of Chicago Press, 2016.
- 8 Reza SM: Transforming big data into computational models for personalized medicine and health care. *Dialogues Clin Neurosci* 2016;18:339–343.
- 9 Safdar S, Zafar S, Zafar N, Khan NF: Machine learning based decision support systems (DSS) for heart disease diagnosis: a review. *Artif Intell Rev* 2018, in press, available online: <https://doi.org/10.1007/s10462-017-9552-8>.
- 10 Pombo N, Araújo P, Viana J: Knowledge discovery in clinical decision support systems for pain management: a systematic review. *Artif Intell Med* 2014;60:1–11.
- 11 Vellido A, Ribas V, Morales C, Ruiz-Santamaría A, Ruiz-Rodríguez J: Machine learning for critical care: state-of-the-art and a sepsis case study. *Biomed Eng Online*, in press.
- 12 Cabitza F, Rasoini R, Gensini GF: Unintended consequences of machine learning in medicine. *JAMA* 2017;318:517–518.
- 13 Goodman B, Flaxman S: European Union regulations on algorithmic decision making and a “right to explanation.” *AI Magazine* 2017;38:50–57.
- 14 Vellido A, Martín-Guerrero JD, Lisboa PJG: Making machine learning models interpretable; in Verleysen M (ed): *Proceedings of the 20th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN 2012)*. Bruges, 2012, pp 163–172.
- 15 Jackups R: Deep learning makes its way to the clinical laboratory. *Clin Chem* 2017;63:1790–1791.

- 16 Tu JV: Advantages and disadvantages of using artificial neural networks versus logistic regression for predicting medical outcomes. *J Clin Epidemiol* 1996;49:1225–1231.
- 17 Reid MJ: Black-box machine learning: implications for healthcare. *Polygeia*, 2017. <https://www.polygeia.com/single-post/2017/04/06/Black-box-machine-learning-implications-for-healthcare>.
- 18 Dominik S, Sedlmair M, Zhang L, Lee JA, Peltonen J, Weiskopf D, North SC, Keim DA: What you see is what you can change: human-centered machine learning by interactive visualization. *Neurocomputing* 2017;268:164–175.
- 19 Bhanot G, Biehl M, Villmann T, Zühlke D: Biomedical data analysis in translational research: integration of expert knowledge and interpretable models; in Verleysen M (ed), 25th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN 2017). Bruges, 2017, pp 177–186.
- 20 Brier ME, Gaweda AE: Artificial intelligence for optimal anemia management in end-stage renal disease. *Kidney Int* 2016;90:259–261.
- 21 Fernández EA, Valtuille R, Balzarini M: Artificial neural networks applications in dialysis; in Azar A (ed): *Modeling and Control of Dialysis Systems*. Studies in Computational Intelligence. Berlin, Springer, 2013, vol 405, pp 1145–1179.
- 22 Chiu JS, Chong CF, Lin YF, Wu CC, Wang YF, Li YC: Applying an artificial neural network to predict total body water in hemodialysis patients. *Am J Nephrol* 2005;25:507–513.
- 23 Wang YF, Hu TM, Wu CC, Yu FC, Fu CM, Lin SH, Huang WH, Chiu JS: Prediction of target range of intact parathyroid hormone in hemodialysis patients with artificial neural network. *Comput Methods Programs Biomed* 2006;83:111–119.
- 24 Martín-Guerrero JD, Gómez F, Soria-Olivas E, Schmidhuber J, Climente-Martí M, Jiménez-Torres NV: A reinforcement learning approach for individualizing erythropoietin dosages in hemodialysis patients. *Expert Syst Appl* 2009;36:9737–9742.
- 25 Barbieri C, Molina M, Ponce P, Tothova M, Cattinelli I, Titapiccolo JJ, Mari F, Amato C, Leipold F, Wehmeyer W, Stuard S: An international observational study suggests that artificial intelligence for clinical decision support optimizes anemia management in hemodialysis patients. *Kidney Int* 2016;90:422–429.
- 26 Rodrigues M, Peixoto H, Esteves M, Machado J: Understanding stroke in dialysis and chronic kidney disease. *Proc Comput Sci* 2017;113:591–596.
- 27 Saadat S, Aziz A, Ahmad H, Imtiaz H, Zara SS, Kazmi A, Aslam S, Naqvi N, Saadat S: Predicting quality of life changes in hemodialysis patients using machine learning: generation of an early warning system. *Cureus* 2017;9:e1713.
- 28 Bobrowski L, Łukaszuk T, Lindholm B, Stenvinkel P, Heimbürger O, Axelsson J, Bárányi P, Carrero JJ, Qureshi AR, Luttropp K, Debowska M, Nordfors L, Schalling M, Waniewski J: Selection of genetic and phenotypic features associated with inflammatory status of patients on dialysis using relaxed linear separability method. *PLoS One* 2014;9:e86630.
- 29 Ricci Z, Romagnoli S, Ronco C: Automatic dialysis and continuous renal replacement therapy: keeping the primacy of human consciousness and fighting the dark side of technology. *Blood Purif* 2017;44:271–275.
- 30 Lankhorst CE, Wish JB: Anemia in renal disease: diagnosis and management. *Blood Rev* 2010;24:39–47.
- 31 Foley RN: Treatment of anemia in chronic kidney disease: known, unknown, and both. *J Blood Med* 2011;2:103–112.
- 32 Martínez-Martínez JM, Escandell-Montero P, Barbieri C, Soria-Olivas E, Mari F, Martínez Sober M, Amato C, Serrano Lopez A, Basi M, Magdalena Benedito R, Stopper A, Gatti E: Prediction of the hemoglobin level in hemodialysis patients using machine learning techniques. *Comput Methods Programs Biomed* 2014;117:208–217.
- 33 Barbieri C, Mari F, Stopper A, Gatti E, Escandell-Montero P, Martínez-Martínez JM, Martín-Guerrero JD: A new machine learning approach for predicting the response to anemia treatment in a large cohort of end stage renal disease patients undergoing dialysis. *Comput Biol Med* 2015;61:56–61.
- 34 Escandell-Montero P, Chermisi M, Martínez-Martínez JM, Gomez Sanchis J, Barbieri C, Soria Olivas E, Mari F, Vila Frances J, Stopper A, Gatti E, Martín Guerrero JD: Optimization of anemia treatment in hemodialysis patients via reinforcement learning. *Artif Intell Med* 2014;62:47–60.
- 35 Stefánsson BV, Brunelli SM, Cabrera C, Rosenbaum D, Anum E, Ramakrishnam K, Jensen DE, Stalhhammar NO: Intradialytic hypotension and risk of cardiovascular disease. *Clin J Am Soc Nephrol* 2014;9:2124–2132.
- 36 Santoro A, Mancini E, Paolini F, Spongano M, Zucchelli P: Automatic control of blood volume trends during hemodialysis. *ASAIO J* 1994;40:M419–M422.
- 37 Leung KCW, Quinn RR, Ravani P, Duff H, MacRae JM: Randomized crossover trial of blood volume monitoring-guided ultrafiltration biofeedback to reduce intradialytic hypotensive episodes with hemodialysis. *Clin J Am Soc Nephrol* 2017;12:1831–1840.
- 38 Nesrallah GE, Suri RS, Thiessen-Philbrook H, Heidenheim P, Lindsay RM: Can extracellular fluid volume expansion in hemodialysis patients be safely reduced using the hemocontrol biofeedback algorithm? A randomized trial. *ASAIO J* 2008;54:270–274.
- 39 Selby NM, Lambie SH, Camici PG, Baker CS, McIntyre CW: Occurrence of regional ventricular dysfunction in patients undergoing standard and biofeedback dialysis. *Am J Kidney Dis* 2006;47:830–841.
- 40 Winkler RE, Grandi F, Santoro A: Blood volume regulation; in Goretti Penido M (ed): *Technical Problems in Patients on Hemodialysis*. London, InTech, 2011, chapt 15.
- 41 Ettema EM, Kuipers J, Groen H, Kema IP, Westerhuis R, de Jong PE, Franssen CFM: Vasopressin release is enhanced by the Hemocontrol biofeedback system and could contribute to better haemodynamic stability during haemodialysis. *Nephrol Dial Transplant* 2012;27:3263–3270.
- 42 Doria M, Genovesi S, Biagi F, Steckiph D, Mancini E, Stella A, Santoro A: The dialysis staff workload and the blood volume tracking system during the hemodialysis sessions of hypotension-prone patients. *Int J Artif Organs* 2014;37:292–298.
- 43 Tijink MS, Wester M, Glorieux G, Gerritsen KG, Sun J, Swart PC, Borneman Z, Wessling M, Vanholder R, Joles JA: Stamatialis: mixed matrix hollow fiber membranes for removal of protein-bound toxins from human plasma. *Biomaterials* 2013;34:7819–7828.
- 44 Agar J: Review: understanding sorbent dialysis systems. *Nephrology* 2010;15:406–411.
- 45 Armignacco P, Garzotto F, Neri M, Lorenzin A, Ronco C: WAK engineering evolution. *Blood Purif* 2015;39:1–3.
- 46 Gura V, Macy AS, Beizai M, Ezon C, Golper TA: Technical breakthroughs in the wearable artificial kidney (WAK). *Clin J Am Soc Nephrol* 2009;4:1441–1448.
- 47 Roy S, Goldman K, Marchant R, Zydny A, Brown D, Fleischman A, Conlisk A, Desai T, Duffy S, Humes H, Fissell W: Implanted renal replacement for end-stage renal disease. *Panminerva Med* 2011;53:155–166.
- 48 Nadkarni GN, Coca SG, Wyatt CM: Big data in nephrology: promises and pitfalls. *Kidney Int* 2016;90:240–241.
- 49 Ketchersid T: Big data in nephrology: friend or foe? *Blood Purif* 2013;36:160–164.
- 50 Russell SJ, Norvig P: *Artificial intelligence: a modern approach*; in Russell SJ, Norvig P (eds): Prentice Hall Series in Artificial Intelligence, ed 3. London, Pearson Education, 2009.
- 51 Goodfellow I, Bengio Y, Courville A: *Deep Learning*. Cambridge, MIT Press, 2016.
- 52 Zeng T, Wu B, Ji S: DeepEM3D: approaching human-level performance on 3D anisotropic EM image segmentation. *Bioinformatics* 2017;33:2555–2562.
- 53 Lake BM, Salakhutdinov R, Tenenbaum JB: Human-level concept learning through probabilistic program induction. *Science* 2015;350:1332–1338.
- 54 Xiong W, Droppo J, Huang X, Seide F, Seltzer M, Stolcke A, Yu D, Zweig G: Achieving human parity in conversational speech recognition. *arXiv* 2016;arXiv:1610.05256v2.

- 55 Yu KH, Zhang Ce, Berry GJ, Altman RB, Ré C, Rubin DL, Snyder M: Predicting non-small cell lung cancer prognosis by fully automated microscopic pathology image features. *Nat Commun* 2016;7:12474.
- 56 Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, Thrun S: Dermatologist-level classification of skin cancer with deep neural networks. *Nature* 2017;542:115–118.
- 57 Weng SF, Reys J, Kai J, Garibaldi JM, Qureshi N: Can machine-learning improve cardiovascular risk prediction using routine clinical data? *PLoS One* 2017;12:0174944.
- 58 How hospitals could be rebuilt, better than before. *The Economist*, April 2017. <https://www.economist.com/news/international/21720278-technology-could-revolutionise-way-they-work-how-hospitals-could-be-rebuilt-better>.
- 59 Bonnefon JF, Shariff A, Rahwan I: The social dilemma of autonomous vehicles. *Science* 2016;352:1573–1576.
- 60 Luetge C: The German ethics code for automated and connected driving. *Philos Technol* 2017;30:547–558.
- 61 Carpenter J: Google’s algorithm shows prestigious job ads to men, but not to women. *Washington Post*, 2015. <https://www.washingtonpost.com/news/the-intersect/wp/2015/07/06/googles-algorithm-shows-prestigious-job-ads-to-men-but-not-to-women-heres-why-that-should-worry-you>.
- 62 Smith J: Crime prediction tool may be reinforcing discriminatory policing. *Business Insider*, 2016. <http://www.businessinsider.com/predictive-policing-discriminatory-police-crime-2016-10>.