© 2015 IMIA and Schattauer GmbH

Health Informatics via Machine Learning for the Clinical Management of Patients

D. A. Clifton, K. E. Niehaus, P. Charlton, G. W. Colopy
Institute of Biomedical Engineering, Department of Engineering Science, University of Oxford, Oxford, UK

Summary

Objectives: To review how health informatics systems based on machine learning methods have impacted the clinical management of patients, by affecting clinical practice.

Methods: We reviewed literature from 2010-2015 from databases such as Pubmed, IEEE xplore, and INSPEC, in which methods based on machine learning are likely to be reported. We bring together a broad body of literature, aiming to identify those leading examples of health informatics that have advanced the methodology of machine learning. While individual methods may have further examples that might be added, we have chosen some of the most representative, informative exemplars in each case.

Results: Our survey highlights that, while much research is taking place in this high-profile field, examples of those that affect the clinical management of patients are seldom found. We show that substantial progress is being made in terms of methodology, often by data scientists working in close collaboration with clinical groups.

Conclusions: Health informatics systems based on machine learning are in their infancy and the translation of such systems into clinical management has yet to be performed at scale.

Keywords

Health informatics, data mining, information systems, electronic health records

Yearb Med Inform 2015;10:38-43
http://dx.doi.org/10.15265/IY-2015-014
Published online August 13, 2015

Abstract

The modern generation of health informatics systems based on machine learning incorporates a range of technologies from (i) wearable sensors for the acquisition of data from patients, through to (ii) biomedical signal processing methods for conditioning the high-resolution data that result, through to (iii) "big data" machine learning approaches to fuse the heterogeneous data that are currently collected from routine clinical care in many healthcare systems. The latter may include fusion of time-series data from sensors; categorical data from electronic health records (EHRs); and, increasingly, biomarkers derived from genomics, proteomics, and other sources. While this field holds substantial promise for the future of medicine, and for our ability to tailor care to the particular physiology of the individual, the penetration of such systems based on machine learning into actual clinical practice is in its infancy. This review surveys the recent literature in this rapidly-changing field, aiming to investigate how health informatics systems that employ machine-learning methods are affecting the clinical management of patients. While the range of literature presented is broad, underpinning each publication is the demonstration of machine learning methods for health informatics.

This review describes health informatics systems that are based on machine learning, throughout the patient journey through a typical hospital healthcare system, from (i) the intensive care unit (ICU), to (ii) discharge and subsequent monitoring on acute wards and on general wards, to (iii) wider-scale tracking of patient condition using the EHR.

1 Clinical Management in the Intensive Care Unit

The ICU routinely cares for the most severely ill patients in the hospital, allowing for the provision of radical lifesaving treatment. Once admitted, the condition of patients is heavily monitored to facilitate the restoration of their physiology to "normality". Many clinically-relevant parameters, such as heart rate (HR) and blood pressure (BP), are monitored as frequently as every second. However, due to the limitation of human resource allocation, much of this information is not used by clinical staff. Furthermore, critical care medicine is one of the fastest-growing medical fields in terms of patient numbers [1], and the complexity of critical illness requires interpretation of numerous interacting parameters. As a result, the ICU has become a promising area for the development of novel informatics for the real-time processing of streams of physiological data, with the aim of providing clinical decision support for patient management.

Prediction of patient outcomes is an immediately obvious application of health informatics and has a long history, which has formed a basis for more recent innovations. In 1981, Knaus et al. published the acute physiology score (APS) as a component of the acute physiology, age, and chronic health evaluation (APACHE) system [2]. With the recent increase in the quantity of available data, and the advances in machine learning techniques that allow these data to be processed, there is a great opportunity for creating accurate patient-specific prediction models, which therefore allow an assessment of individual patients, suitable for affecting management decisions. Recent work

by Johnson et al. applied a hybrid system combining a genetic algorithm with particle swarm optimisation to create an ICU severity score that used a parsimoniously-selected feature set [3]. However, these models use aggregate summaries of patient physiology over the first hour or day in the ICU (such as the maximum HR). Constraining the analysis to considering only early-admission data in this way is, as the authors describe, due to the traditional setting of ICU severity scores for use at admission.

A recent review of all mortality prediction studies found 94 investigations that contained 118 predictive tools [4]. Of these 118 tools, 34 were unique. This demonstrated that there is a lack of wide uptake of more sophisticated mortality prediction models. However, there exist promising approaches undertaken in recent years. Saria et al. present a model which aimed to predict neonate morbidity using continuous vital-sign monitoring data from the first 3 hours of the neonate's life [5]. The model first filtered the continuous vital-sign data into long- and short-term variability signals. For each class of patient, a probability distribution for each vital sign was estimated, and the probabilities of morbidity from each of these distributions were aggregated into a single score using regularised logistic regression. The resulting score, named the PhysiScore, had excellent discrimination for various neonatal morbidities as compared to more commonly utilized severity scores. To ascertain whether the dynamics of the heart rate could further inform a predictive model, the same authors then apply an extension of Latent Dirichlet Allocation (LDA) for time-series data. The model creates "topics" for each time series which are shared across patients. These topics refer to concepts such as "healthy" or "lung", and control the distribution of autoregressive models which describe the data. The result was validated with 30 infants and demonstrated that the length of time a neonate's conditioned was deemed to spend in the "healthy" topic was correlated with severity of that neonate's condition.

Similar models have been used to monitor adult ICU patients. Lehman et al. apply a switching vector autoregressive (SVAR) model to the dynamics of HR and BP time-series for patients in the ICU [6]. This model similarly summarises the time-series using autoregressive modes and uses the proportion of time in each mode as a marker of health or severity. The authors showed that BP dynamics alone are able to outperform existing ICU scores. Furthermore, they show that the addition of BP dynamics to the widely-used APACHE IV score further improved its AUROC from 0.82 to 0.85.

One particularly successful health informatics system is the *HeRO* monitor, which has been validated via a randomised controlled trial [7]. This system provides early prediction of neonatal sepsis by incorporating various measures of HR variability including the standard deviation of R-R Intervals, sample asymmetry, and sample entropy. These metrics captured a common pattern in HR observed in neonates with sepsis: rapid decelerations combined with low variability. A comparison of the control and study population demonstrated that for every 48 low birth weight infants, the HeRO monitor saved one life.

In addition to the numerous time-series signals acquired in the ICU, there also exists a large body of useful information contained within the clinical notes written during a patient's stay in the ICU. These notes usually contain information regarding past medical history, diagnostic information, current patient status, treatment plans and discussion of end-of-life care options. These notes represent an underused, but highly relevant, source of information, and are now often recorded electronically, making them more suitable for automated interpretation. Ghassemi et al. demonstrated the utility of incorporating information contained in these notes into ICU-based systems [8]. The system extracts a set of 50 topics from the case notes for patients, using an LDA model derived from the top 500 words from the notes. The ability of the model to predict long-term mortality increased AUROC from 0.711 to 0.788 when such data were included in the ICU's severity score. Further research by the same authors demonstrated that similar performance could be achieved by an algorithm implemented at the bedside [9]. The AUROC of the latter model for predicting mortality (without information from the ICU notes) was 0.76, while the addition of learnt topics from the notes improved this to 0.84.

2 Clinical Management Outside the ICU: Wearable Sensors

The majority of patients in hospitals are typically ambulatory, and in this setting patient condition is often monitored only by manual observation by the clinical staff. It is established that the use of wearable sensors for managing the care of such populations can improve patient outcomes. A recent study by Cleverley et al. [10] have found that remote monitoring is associated with improved survival to hospital discharge, following in-hospital cardiac arrest. However, despite such studies, evidence suggests that the current use of systems based on wearable sensors in hospital is far from ideal [11]. This has been attributed to systems being used inappropriately, in an attempt to manage patients for whom such monitoring confers no benefit [12,13]. This latter effect has been attributed to the inadequacy of current guidelines for stratifying patients to determine whether or not additional monitoring will be of benefit [14].

The standard of clinical management for patients on acute hospital wards is typically to have vital signs measured manually every 4-6 hours [15]. There is evidence that wearable sensors are being used in clinical practice to increase the frequency of measurement of the vital signs, with the aim of identifying the deteriorating patient [16-18]. Automated analysis of the trends of vital signs acquired from wearable sensors has been used in the care of ambulatory post-operative patients with the aim of determining if the time-series data from a patient appear similar to those of "normal" or "at risk" patients [19-21]. This latter study used Bayesian Gaussian processes (GPs) to perform clustering of entire univariate time-series; a similar approach was described by Duerichen et al. [22], in which multiple time-series of vital signs were modelled simultaneously using multi-task Gaussian processes, for the monitoring of ambulatory patients. Vital-sign data from wearable sensors were fused with the manual observations made by the clinical staff for acutely ill, ambulatory patients in Clifton et al. [23-

25], using GPs and one-class support vector machines. Such methods have been extended by the use of health informatics based on extreme value theory, a branch of statistics used to model extremal data, and which have been demonstrated in a patient monitoring study in a US hospital [26, 27]. The evidence base for the use of health informatics systems based on machine learning was questioned in [28], where methods for evaluating the efficacy of a system used for identifying the deteriorating patient in a large Emergency Department setting were described.

Gaussian processes are non-parametric Bayesian methods of performing time-series analysis, which are becoming increasingly adopted by probabilistic health informatics systems based on machine learning. They offer a principled manner of coping with noisy or incomplete data, as typically encountered in routine clinical management. We note in passing that further applications of this approach outside the realm of sensor-based analysis include the assessment of heat-exposure sickness [29, 30], malaria infections [31], and the management of patient billing records [32].

Other recent studies, such as [33], focussed on using existing sensors (such as ECG electrodes) for the assessment of respiratory rate (RR) from cohorts of similarly ambulatory patients. This latter area of application has also been addressed by the use of pulse oximeters, again worn by ambulatory patients [34-36]. The penetration of such systems into actual clinical management is limited, and this is a point addressed by Orphanidou et al. [37] in which algorithms were developed to improve both the robustness of RR estimation from pulse oximeters, and to extend their battery life such that they could be used in the clinic without frequent replacement of batteries. The use of pulse oximeters to estimate other vital signs, such as changes in BP from the pulse transit time (the interval between a QRS complex from the ECG and the corresponding pulse of a finger-worn oximeter), has been well-described in the literature, but a recent study [38] claims to have improved the perceived lack of robustness that has previously prevented the use of such methods in clinical practice.

3 Clinical Management Outside the ICU, via "Big Data" Health Informatics

Beyond the traditional monitoring of vital signs, there is a rich resource of data collected from the routine care of patients in many hospitals. This "big data" problem typically means that only specific subsets of data are used for clinical management, whereas health informatics systems based on machine learning offer the potential for informing clinical care by fusing information from the EHR. This section considers those EHR-based analyses that pertain to applications outside the ICU (where the latter were described in section 1).

Two prominent themes within health informatics concern the integration of genomic data into patient care and the intelligent use (and re-use) of EHR data [39-41]. While research output in these areas is increasing, translation into routine clinical practice is still slow. As Bright et al. report in their review of clinical decision-support systems, while there are many small-scale examples of local adoption of clinical support tools, there is still a dearth of evidence on multi-site clinical outcomes [42]. There are also few examples of predictive systems, even in the research stage, that employ both clinical and genomic features. Partially, this is due to our still rudimentary understanding of human genomics: in 2014, Dewey et al. reported on their experience of whole-genome sequencing (WGS) twelve healthy subjects [43]. The predominant finding was the WGS is not yet ready for use in clinical practice, given the inconsistency of results, inadequate base overage in key disease-related areas of the genome, and rudimentary understanding of how the majority of SNPs may affect the risk of certain diseases. Much informatics research, therefore, is focused upon developing tools that will enable us to answer more fundamental questions of human physiology. EHR phenotyping algorithms, which are designed to enable large-scale genome-wide association studies (GWAS), fall into this category. Though this research is not intended to produce immediate clinical decision support tools, it will lead to the development of appropriate therapies and diagnostic tools further down the line.

The EHR is a rich source of data for many avenues of research outside the ICU, providing input data for both traditional epidemiology studies and training data for the development of predictive algorithms. However, patient records are often partially missing, incorrect, systematically misleading, and contradictory. Particularly relevant for supervised machine learning analysis, patient labels may be confounded by the fact that care is being provided [44]. Hripcsak and Albers [45] outline several of the associated challenged of extracting EHR data that reflects true patient physiology.

The Electronic Medical Records and Genomics Network (eMERGE) is a large project funded by the National Human Genome Research Institute in the US, with the goal of hastening the use of genome sequencing in clinical practice. The project's first phase (eMERGE-1) formed a consortium of five large clinical sites, each with access to both EHR and biobank data. The primary goal was the establishment of EHR-phenotyping algorithms, to enable GWAS by linking large EHR-phenotyped patient cohorts with their genomic biobank data [46]. The eMERGE-1 project was quite successful in extracting phenotypes from the EHR, publishing algorithms for phenotypes ranging from rheumatoid arthritis to drug-induced liver disease. These algorithms are publicly-available and maintained at the Phenotype KnowledgeBase (PheKB). In the second phase of eMERGE funding (eMERGE-II), the focus has shifted towards implementing genomic tools as clinical support systems in EHRs. The *PGx* project is the first of these translational informatics projects within the eMERGE network, and aims to implement a system that will use knowledge of pharmacogenomic relationships to alert physicians when a patient may require an adjusted medication or dosage [47].

This research is not confined to the eMERGE network; for example, Wang, et al. used semi-supervised machine learning to extract ovarian cancer diagnoses from the EHR [48]. Shivade et al. recently conducted a review of phenotyping algorithms, finding 97 studies published between 2010-2013. Many of these have used natural-language processing (NLP) and rule-based defini-

tions in their phenotype definitions, often followed by machine learning or statistical analyses on the resulting phenotype-extracted patients' data [49].

The integrating biology and the bedside (i2b2) initiative [50], the UK Biobank, and the Kaiser Permanent Research Program on Genes, Environment, and Health (RPGEH) also all similarly aim to combine large amounts of clinical and genomic data; Jensen et al. provide additional examples [40]. The EHR-based phenotyping efforts described above have largely been focused upon identifying patients who can be described by current disease terminology. However, with growing access to EHR data, it is now becoming possible to identify novel patient subgroups. While many complex diseases have traditionally been thought of as monolithic entities, it is very likely that finer-grained phenotypic sub-types exist [51,52]. Ho et al. have framed the problem in terms of tensors, using nonnegative tensor factorisation to more automatically identify patient phenotypes [53].

Doshi-Velez, et al., for example, looked at the timing and presence of ICD9 diagnostic codes in the EHRs of over 13,000 autism spectrum disorder patients [54]. Using hierarchical clustering, they found three distinct patient trajectories, corresponding to a group of patients characterised by seizures, a group by multisystem disorders, and a group by psychiatric disorders. Schulam, et al. built upon such approaches by developing a probabilistic model to cluster patient trajectories of various clinical indicators related to scleroderma [55]. Lasko et al. used Gaussian process regression over serum uric acid measurements, combined with a two-layered neural network, to create a sparse feature set of time-series trajectories. When projected into two dimensions, the learned feature sets showed good discrimination between known diseases and also suggested possible additional underlying disease subtypes [56]. The goal of such research is primarily hypothesis-generating, with the aim of determining if patients within phenotypic subgroups are united by a common genomic pathway, or if they might respond to certain types of treatment. Such analysis is not intended to be immediately translatable into the clinic, but informs our clinical understanding of these disorders and points towards promising future areas of research.

While the studies presented above enable the study of physiological mechanisms that may be involved in human disease, there are many examples of recently-developed systems that are designed to be incorporated into an intelligent EHR system. Huang et al. used features derived from the EHR to create a screening tool for depression with sensitivities and specificities on par with those achieved by primary care physicians [57]. Wiens et al. created a risk-stratification model for becoming infected with *Clostridium difficile* by incorporating the time-based relationship between clinical variables in the EHR; the resulting system was able to predict patient risk much more successfully than when considering features in aggregate [58]. Van der Heijden et al. used temporal Bayesian networks to predict exacerbation events for chronic obstructive pulmonary disease patients based on EHR data; their system was able to predict an exacerbation within 24 hours with an AUROC of 0.90 in a validation study [59].

On a larger scale, Li et al. illustrate how EHR data can be mined and combined with curated genomic data to produce early prognostic risk indicators for common medical conditions [60]. The goal of such studies is for the resulting tools to be adopted into clinical practice as new screening techniques. One example of the successful translation of a genomics-based informatics tool has been the development of gene-expression based algorithms for breast cancer prognosis. Van De Vijver et al. [61] used supervised and unsupervised machine learning techniques on gene expression data from breast cancer cells to provide large improvement over existing clinical prognostic markers. These techniques have now been developed into a commercial tool (MammaPrint, Agendia, Netherlands) that is used for breast cancer prognosis in the clinic. The visualisation of such complex data was addressed by Ding et al., who developed a visualisation platform using unsupervised learning methods [62]. Das et al. provide a review of such informatics techniques and data types that have been used for predicting cancer prognoses [63], with the aim of improving the clinical management and treatment of such patients.

Clinical microbiology is another field that is set to be transformed through genome sequencing and supportive informatics tools embedded within the EHR [64]. While genome sequencing, combined with machine learning algorithms, has been used for HIV drug resistance prediction for many years, similar methodologies are now being applied for bacterial infections. For instance, Stoesser et al. [65] and Gordon et al. [66] have illustrated the feasibility of using whole-genome sequences from *Staphylococcus aureus*, *Klebsiella pneumonia*, and *Escherichia coli* to determine bacterial drug resistance. Niehaus et al. illustrated the first steps towards creation of a machine-learning-based antibiotic resistance-prediction tool with the aim of providing appropriate treatments [67].

4 Conclusions

We conclude by emphasising that the field of health informatics systems based on machine learning, drawing on disparate datatypes from the ICU, the wider hospital, and from (potentially very complex) EHR data, is in its infancy. While the majority of hospitals in the developed world have implemented EHR systems of some kind, the integrated use of the large quantities of data that arise from such systems is not employed at scale. This article has surveyed recent developments in this field, in which the clinical management of patients has been affected by health informatics systems based on machine learning, or in which systems for performing such management are in development. We observe in closing that the barriers-to-entry for such activity are substantial, and are affected by the availability of multidisciplinary teams drawing on both clinical expertise and practitioners from the information sciences. The implications for research in this field are that clinical management of patients will become more dependent on machine-learning systems with the ever-increasing quantities of data that are being collected in routine care, from the ICU through to discharge, in most healthcare systems.

References

1. Vincent JL. Critical care: where have we been and where are we going? Crit Care 2013;17(1):S2.
2. Knaus WA, Zimmerman JE, Wagner DP, Draper EA, Lawrence DE. APACHE, acute physiology and chronic health evaluation: a physiologically based classification system. Crit Care Med 1981;9(8):591-7.
3. Johnson AEW, Kramer AE, Clifford GD. A new severity of illness scale using a subset of acute physiology and chronic health evaluation data elements shows comparable predictive accuracy. Crit Care Med 2013;41(7):1711-8.
4. Siontis GCM, Tzoulaki I, Ioannidis JPA. Predicting death: an empirical evaluation of predictive tools for mortality. Arch Intern Med 2011;171(19):1721-6.
5. Saria S, Rajani AK, Gould J, Koller D, Penn AA. Integration of early physiological responses predicts later illness severity in preterm infants. Sci Transl Med 2010;2(48):48-65.
6. Lehman LW, Adams RP, Mayaud L, Moody GB, Malhotra A, Mark RG. A physiological time series dynamics-based approach to patient monitoring and outcome prediction. IEEE J Biomed Health Inform 2015;19(3):1068-76.
7. Moorman JR, Carlo WA, Kattwinkel J, Schelonka RL, Porcelli PJ, Navarrete CT, et al. Mortality reduction by heart rate characteristic monitoring in very low birth weight neonates: a randomized trial. J Pediatr 2011;159(6): 900-6.
8. Ghassemi M, Naumann T, Joshi R, Rumshisky A. Topic models for mortality modelling in intensive care units. Proceedings of International Conference on Machine Learning; 2012. p. 1-4.
9. Ghassemi M, Naumann, T, Doshi-Velez F, Brimmer N, Joshi R, Rumshisky A, et al. Unfolding physiological state: mortality modelling in intensive care units. Proceedings of Knowledge Discovery and Data Mining; 2014. p. 75-84.
10. Cleverley K, Mousavi N, Stronger L, Ann-Bordun K, Hall L, Tam JW, et al. The impact of telemetry on survival of in-hospital cardiac arrests in non-critical care patients. Resuscitation 2013;84(7):878–82.
11. Chen EH. Appropriate use of telemetry monitoring in hospitalized patients. Current Emergency Hospital Medical Reports 2013;2(1):52–6.
12. Benjamin E, Klugman RA, Luckmann R, Fairchild DG, Abookire SA. Impact of cardiac telemetry on patient safety and cost. Am J Manag Care 2013;19(6):e225–232.
13. Bulger J, Nickel W, Messler J, Goldstein J, O'Callaghan J, Auron M, et al. Choosing wisely in adult hospital medicine: five opportunities for improved healthcare value. J Hosp Med 2013;8(9):486–92.
14. Falun N, Nordrehaug JE, Hoff PI, Langorgen J, Moons P, Norekval TM. Evaluation of the appropriateness and outcome of in-hospital telemetry monitoring. Am J Cardiol 2013;112(8):1219–23.
15. Walsh JA, Topol EJ, Steinhubl SR. Novel wireless devices for cardiac monitoring. Circulation 2014;130(7):573–81.
16. Najafi N, Auerbach A. Use and outcomes of telemetry monitoring on a medicine service. Arch Intern Med 2012;172(17):1349-50.
17. Bonnici T, Charlton P, Alastruey J, Tarassenko L, Watkinson P, Beale R. Continuous physiological monitoring of ambulatory patients. Proceedings of MECBioeng; 2014: 38.
18. Welch J, Moon J, McCombie S. Early detection of the deteriorating patient: the case for a multi-parameter patient-worn monitor. Biomed Instrum Technol 2012;46(s2):57–64.
19. Pimentel MAF, Clifton DA, Tarassenko L. Gaussian process clustering for the functional characterisation of vital-sign trajectories. Proceedings of IEEE Machine Learning for Signal Processing 2013: 1–6.
20. Clifton DA, Clifton L, Hugueny S, Wong D, Tarassenko L. An extreme function theory for novelty detection. IEEE J Sel Top Signal Process 2013;7(1):28-37.
21. Pimentel MAF, Clifton DA, Clifton L, Watkinson PJ, Tarassenko L. Modelling physiological deterioration in post-operative patient vital-sign data. Med Biol Eng Comput 2013;51:869-77.
22. Duerichen, R, Pimentel MAF, Clifton L, Schweikard A, Clifton DA. Multi-task Gaussian processes for multivariate physiological time-series analysis. IEEE Trans Biomed Eng 2015 Jan;62(1):314-22.
23. Clifton L, Clifton DA, Pimentel MAF, Watkinson PJ, Tarassenko L. Predictive monitoring of mobile patients by combining clinical observations with data from wearable sensors. IEEE J Biomed Health Inform 2014;18(3):722-30.
24. Clifton L, Clifton DA, Zhang Y, Watkinson PJ, Tarassenko L, Yin, H. Probabilistic novelty detection with support vector machines. IEEE Transactions on Reliability 2014;63(2):455-67.
25. Clifton L, Clifton DA, Pimentel MAF, Watkinson PJ, Tarassenko L. Gaussian processes for personalised e-health monitoring with wearable sensor. IEEE Trans Biomed Eng 2013;60(1):193-7.
26. Clifton DA, Hugueny S, Clifton L, Tarassenko L. Extending the generalised Pareto distribution for novelty detection in high-dimensional spaces. J Signal Process Syst 2014;74:323-39.
27. Clifton DA, Hugueny S, Tarassenko L. Novelty detection with multivariate extreme value statistics. J Signal Process Syst 2011;65:371-89.
28. Clifton DA, Wong D, Clifton L, Pullinger R, Tarassenko L. A large-scale clinical validation of an integrated monitoring system in the Emergency Department. IEEE J Biomed Health Inform 2013;17(4):835-42.
29. Heaton MJ, Peng RD. Flexible distributed lag models using random functions with application to estimating mortality displacement from heat-related deaths. J Agric Biol Environ Stat 2012;17(3):313-31.
30. Heaton MJ, Sain SR, Monaghan AJ, Wilhelmi OV, Hayden MH. An analysis of an incomplete marked point pattern of heat-related 911 calls. J Am Stat Assoc 2015;110(509):123-35.
31. Lloyd C, Gunter T, Osborne MA, Roberts SJ. Variational inference for Gaussian process modulated Poisson processes. Proceedings of International Conference on Machine Learning; 2015. p. 1-9.
32. Lasko TA. Efficient inference of Gaussian process modulated renewal processes with application to medical event data. Proceedings of Uncertainty in Artificial Intelligence; 2014. p. 36.
33. Donnelly N, Hunniford T, Harper R, Flynn A, Kennedy A, Branagh D, et al. Demonstrating the accuracy of an in-hospital ambulatory patient monitoring solution in measuring respiratory rate. Conf Proc IEEE Eng Med Biol Soc 2013:6711–5.
34. OBS Medical Ltd. Visensia Mobile and its innovative software for automated respiration rate calculation. Available: http://www.obsmedical.com/news/article/visensia-mobile-and-its-innovative-software-for-automated-respiration-rate. Accessed: 30-Nov-2014.
35. Pimentel MAF, Charlton PH, Clifton DA. Probabilistic estimation of respiratory rate from wearable sensors. In: Wearable Electronic Sensors. Mukhopadhyay S, editor. Smart Sensors, Measurement and Instrumentation 2015;15: 241-62.
36. Pimentel MAF, Clifton DA, Clifton L, Tarassenko L. Probabilistic estimation of respiratory rate using Gaussian processes. Conf Proc IEEE Eng Med Biol Soc 2013:2902-5.
37. Orphanidou C, Bonnici T, Charlton P, Clifton DA, Vallance, D, Tarassenko L. Signal quality indices for the electrocardiogram and photoplethysmogram: derivation and applications to wireless monitoring. IEEE J Biomed Health Inform 2015;19(3):832-8.
38. Sempeles S. Continuous wireless monitoring device passes test in first hospital use. J Clin Eng 2013;38(3):86–7.
39. Denny JC. Mining electronic health records in the genomics era. PLoS Comput Biol 2012;8(12):e1002823.
40. Jensen PB, Jensen LJ, Brunak. Mining electronic health records: towards better research applications and clinical care. Nat Rev Genet 2012;13(6):395-405.
41. Simpao AF, Ahumada LM, Galvez JA, Rehman MA. A review of analytics and clinical informatics in health care. J Med Syst 2014;38(4):45.
42. Bright TJ, Wong A, Dhurjati R, Bristow E, Bastian L, Coeytaux RR, et al. Effect of clinical decision-support systems: a systematic review. Ann Intern Med 2012;157(1): 29-43.
43. Dewey FE, Grove ME, Pan C, Goldstein BA, Bernstein JA, Chaib H, et al. Clinical interpretation and implications of whole-genome sequencing. J Am Med Assoc 2014;311(10):1035-45.
44. Paxton C, Niculescu-Mizil A, Saria S. Developing predictive models using electronic medical records: challenges and pitfalls. AMIA Annu Symp Proc 2013:1109-15.
45. Hripcsak G, Albers DJ. Next-generation phenotyping of electronic health records. J Am Med Inform Assoc 2013;20(1):117-21.
46. Gottesman O, Kuivaniemi H, Tromp G, Faucett WA, Li R, Manolio TA, et al. The Electronic Medical Records and Genomics (eMERGE) Network: past, present, and future. Genet Med 2013;15(10):761-71.
47. Rasmussen-Torvik LJ, Stallings SC, Gordon AS, Almoguera B, Basford MA, Bielinski SJ, et al. Design and anticipated outcomes of the eMERGE-PGx project: a multicenter pilot for preemptive pharmacogenomics in electronic health record systems. Clin Pharmacol Ther 2014;96(4):482-9.
48. Wang Z, Shah AD, Tate AR, Denaxas S, Shawe-Taylor J, Hemingway H. Extracting diag-

noses and investigation results from unstructured text in electronic health records by semi-supervised machine learning. PLoS One 2012;7(1):e30412.
49. Shivade C, Raghavan P, Fosler-Lussier E, Embi PJ, Elhadad N, Honson SB, et al. A review of approaches to identifying patient phenotype cohorts using electronic health records. J Am Med Inform Assoc 2014;21(2):221-30.
50. Kohane IS, Churchill SE, Murphy SN. A translational engine at the national scale: informatics for integrating biology and the bedside. J Am Med Inform Assoc 2012;19(2):181-5.
51. Kohane I. Deeper, longer phenotyping to accelerate the discovery of the genetic architectures of diseases. Genome Biol 2014;15(5):115.
52. Loscalzo J, Kohane I, Barabasi AL. Human disease classification in the postgenomic era: a complex systems approach to human pathobiology. Mol Syst Biol 2007;3:124.
53. Ho JC, Ghosh J, Steinhubl S, Stewart WF, Denny JC Malin BA Sun J. Limestone: High-throughput candidate phenotype generation via tensor factorization. J Biomed Inform 2014; 52:199-211.
54. Doshi-Velez, F, Ge Y, Kohane I. Comorbidity clusters in autism spectrum disorders: an electronic health record time-series analysis. Pediatrics 2014;133(1):e54-63.
55. Schulam P, Wigley F, Saria S. Clustering longitudinal clinical marker trajectories from electronic health data: applications to phenotyping and endotype discovery. Proceedings of AAAI Conference on Artificial Intelligence; 2015. p. 1-9.
56. Lasko TA, Denn JC, Levy MA. Computational phenotype discovery using unsupervised feature learning over noisy, sparse, and irregular clinical data. PloS One 2013:8(6):e66341.
57. Huang SH, LePendu P, Iyer SV, Tai-Seale M, Carrell D, Shah NH. Toward personalizing treatment for depression: predicting diagnosis and severity. J Med Inform Assoc 2014;21(6):1069-75.
58. Wiens J, Guttag J, Horvitz E. Patient risk stratification for hospital-associated C. diff as a time-series classification task. Proceedings of Neural Information Processing Systems; 2012. p. 476-84.
59. van der Heijden M, Velikova M, Lucas PJ. Learning Bayesian networks for clinical time series analysis. J Biomed Inform 2014;48:94-105.
60. Li L, Ruau DJ, Patel CJ, Weber SC, Chen R, Tatonetti NP, et al. Disease risk factors identified through shared genetic architecture and electronic medical records. Sci Transl Med 2014;6(234):234ra57.
61. Van de Vijver MJ, He YD, van 't Veer L. A gene-expression signature as a predictor of survival in breast cancer. N Engl J Med 2002;347(25):1999-2009.
62. Ding H, Wang C, Huang K, Machiraju R. iGPSe: a visual analytic system for integrative genomic based cancer patient stratification. BMC Bioinformatics 2014;15:203.
63. Das J, Gayvert KM, Yu H. Predicting cancer prognosis using functional genomics data sets. Cancer Inform 2014;13(S5):85-8.
64. Didelot X, Bowden R, Wilson DJ, Peto TE, Crook DW. Transforming clinical microbiology with bacterial genome sequencing. Nat Rev Genet 2012;13(9):601-12.
65. Stoesser N, Batty EM, Eyre DW, Morgan M, Wyllie DH, Del Ojo Elias C, et al. Predicting antimicrobial susceptibilities for Escherichia coli and Klebsiella pneumoniae isolates using whole genomic sequence data. J Antimicrob Chemother 2013;68(10):2234-44.
66. Gordon NC, Price JR, Cole K, Everitt R, Morgan M, Finney J, et al. Prediction of Staphylococcus aureus antimicrobial resistance by whole-genome sequencing. J Clin Microbiol 2014;52(4):1182-91.
67. Niehaus KE, Walker TM, Crook DW, Peto TE, Clifton DA. Machine learning for the prediction of antibacterial susceptibility in Mycobacterium tuberculosis. Proceedings of IEEE Biomedical Health Informatics; 2014. p. 618-21.

Correspndence to:
David A. Clifton
Institute of Biomedical Engineering
Department of Engineering Science
University of Oxford
Oxford, UK
E-mail: davidc@robots.ox.ac.uk