CLINICAL RESEARCH ARTICLE

Modelling outcomes after paediatric brain injury with admission laboratory values: a machine-learning approach

This article has been corrected since Advance Online Publication and a correction is also printed in this issue

Saeed Kayhanian^{1,2}, Adam M. H. Young¹, Chaitanya Mangla^{2,3}, Ibrahim Jalloh¹, Helen M. Fernandes¹, Matthew R. Garnett¹, Peter J. Hutchinson¹ and Shruti Agrawal⁴

BACKGROUND: Severe traumatic brain injury (TBI) is a leading cause of mortality in children, but the accurate prediction of outcomes at the point of admission remains very challenging. Admission laboratory results are a promising potential source of prognostic data, but have not been widely explored in paediatric cohorts. Herein, we use machine-learning methods to analyse 14 different serum parameters together and develop a prognostic model to predict 6-month outcomes in children with severe TBI.

METHODS: A retrospective review of patients admitted to Cambridge University Hospital's Paediatric Intensive Care Unit between 2009 and 2013 with a TBI. The data for 14 admission serum parameters were recorded. Logistic regression and a support vector machine (SVM) were trained with these data against dichotomised outcomes from the recorded 6-month Glasgow Outcome Scale.

RESULTS: Ninety-four patients were identified. Admission levels of lactate, H+, and glucose were identified as being the most informative of 6-month outcomes. Four different models were produced. The SVM using just the three most informative parameters was the best able to predict favourable outcomes at 6 months (sensitivity = 80%, specificity = 99%).

CONCLUSIONS: Our results demonstrate the potential for highly accurate outcome prediction after severe paediatric TBI using admission laboratory data.

Pediatric Research (2019) 86:641–645; <https://doi.org/10.1038/s41390-019-0510-9>

INTRODUCTION

Traumatic brain injury (TBI) remains a worldwide public health concern, carrying a significant burden of mortality and morbidity. TBI is a particular burden in the paediatric population, where it is estimated to be the leading cause of death in children over the age of 1 year.^{1,2}

An accurate prognostic model in the context of TBI is valuable for both clinical and research purposes. First, it aids clinical decision-making about a patient's monitoring requirements and the necessity of surgical interventions. The extent of uncertainty in making these decisions has been recently highlighted by the results of the BEST-TRIP (Benchmark Evidence from South American Trials: Treatment of Intracranial Pressure) randomised control trial (RCT), which has challenged the wisdom of implementing invasive monitoring for all severe TBI patients, and suggests that much better stratification of patients with this presentation may be required.³ Prognostic models can also be used to inform and direct research about the disease process and the risk factors for progression of secondary injury after TBI.

A number of models using radiological and clinical parameters, such as initial computed tomography (CT) scan and Glasgow Coma Score, to predict outcomes, have been explored and validated in large adult cohorts, but the evidence for the validity of these models in children is less compelling.^{4,5} Admission laboratory

variables have also been explored in adult studies for prognostication, including a study on patients from the large IMPACT database.⁶ However, again, studies of these markers for paediatric cohorts are scant. This is despite the fact that children are known to differ in their pathophysiology after TBI and can generally expect better outcomes compared with similar injury in adults, necessitating specific paediatric evidence.⁷

Admission variables are a particularly useful potential source of prognostication: they are easily and routinely available (even in low- and middle-income countries [LMICs], where CT scans may not be), they may provide a possible avenue of clinical intervention (since many of the variables can be corrected), and deranged values that are highly predictive may inform research efforts into the pathophysiology underlying brain injuries, which are currently poorly understood.

Previous studies examining these variables (including a small number of paediatric studies) have focused on linear and, occasionally, non-linear models for single variables. However, given the numerous laboratory variables available and the possibility for interdependence and cross-correlation, using machine-learning methods to analyse all variables together in various permutations may help to provide the most robust interrogation of the data.

In this context, we conducted this study to identify which, if any, admission laboratory variables are correlated to outcomes after TBI

¹Department of Clinical Neurosciences, Division of Neurosurgery, Addenbrooke's Hospital, University of Cambridge, Cambridge, UK; ²Fitzwilliam College, University of Cambridge, Cambridge, UK; ³Department of Computer Science and Technology, University of Cambridge, Cambridge, UK and ⁴Department of Paediatric Intensive Care, Addenbrooke's Hospital, University of Cambridge, Cambridge, UK

Correspondence: Saeed Kayhanian (sk776@cam.ac.uk)

These authors contributed equally: Saeed Kayhanian, Adam M. H. Young

Received: 19 March 2019 Revised: 17 June 2019 Accepted: 4 July 2019

Published online: 26 July 2019

in children and to explore prediction of outcomes, using both univariate analysis and supervised learning methods.

METHODS

Patient selection and data collection

The clinical records of patients admitted to Cambridge University Hospital's (CUH) Paediatric Intensive Care Unit (PICU) between January 2009 and December 2013 were retrospectively reviewed.

Patients aged 16 years and under, who had sustained a severe TBI, were included for review. All patients received full active management of their TBI according to international guidelines, with a tiered treatment protocol based on positioning, sedation, muscle paralysis, moderate hyperventilation, ventriculostomy, osmotic agents and induced hypothermia, aimed at maintaining intracranial pressure below 20 mmHg.^{8,9} Inclusion criteria for the study were as follows: (1) confirmed TBI, using CT or magnetic resonance imaging; (2) severe injury necessitating continued admission in PICU after day 1 post injury, due to continued requirement for sedation and/or close neurological monitoring; and (3) patients requiring invasive monitoring of intracranial pressure or arterial blood pressure.

The electronic clinical records were examined for the data of laboratory blood tests taken in the Emergency Department on admission. Data for 14 serum parameters were recorded, namely: glucose, haemoglobin, albumin, C-reactive protein, sodium, urea, magnesium, lactate, venous pH, white cell count (total), neutrophil count, haematocrit, prothrombin time and activated partial thromboplastin time.

All patients had been followed up at 6 months and the Glasgow Outcome Scale (GOS) [dead = GOS 1, vegetative state = GOS 2, severe disability = GOS 3, moderate disability = GOS 4, good recovery = GOS 5] was assessed 6 months after initial presentation in the clinic and recorded in the electronic clinical records. The assessments were made based on age-appropriate markers for neurological function and activities of daily living.

Statistical methods

The GOS for each patient was dichotomised as favourable (GOS 4–5) vs. unfavourable (GOS 1–3). Two machine-learning models were trained using these data, namely, linear regression and support vector machine (SVM). The models were evaluated using five-fold cross-validation by averaging the relevant metrics across the folds. This cross-validated metric more accurately reflects the expected performance of our models for new test data that may be different from any previously observed training data. In the case of SVM models, the model parameters are optimised to maximise this cross-validated metric, which prevents our selected models from over-fitting to the training data and consequently ensures our predictions generalise to new test data. Statistical analysis was carried out in Python, using the Scikit-learn and minipy libraries, and *P* values were calculated using the two-tailed Mann–Whitney *U* test. Plots were produced using the matplotlib Python library and the ggplot2 R library.

Inspection of variables. The mutual information (MI) between the blood result variables and the bipartite outcome was examined, along with the correlation coefficient (CC) of the same relationship. Higher MI between a variable and the outcome indicates, broadly speaking, that the variable is more informative of the outcome. MI was measured using the maximal information coefficient (MIC). MIC lies in the range [0, 1], where 0 implies that the variable is uninformative of the outcome, and 1 represents a noise-free relationship between the variable and the outcome, wherein the relationship may be of any form, linear or non-linear. The CC on the other hand is only an indication of the linear relationship between the variables and the outcome, and we examine the absolute value of the CC such that it falls in the same range of [0, 1] as the MIC, with similar interpretation.

Higher absolute CC implies higher MIC, but a highly informative variable, which will have a high MIC, may have a low absolute CC if it has a highly non-linear relationship with the outcome.

Logistic regression. A logistic regression model was trained and then evaluated by constructing receiver-operating characteristic (ROC) curves of blood parameters against bipartite outcomes for each fold during the five-fold cross-validation to then produce an averaged ROC curve. The area under the curve (AUC) was then calculated, with AUC >0.8 generally considered to indicate good performance. Youden's *J*-statistic was used to calculate the optimal decision threshold, balancing sensitivity and specificity for predicting good outcomes.

Support vector machine. SVMs are supervised learning classifiers with state-of-the-art prediction performance. When used with the radial basis function (RBF) kernel, they are capable of learning complex non-linear relationships between the variables and the classes. SVM-based models using RBF kernels were developed in this work for predicting the bipartite outcome from serum result variables. The models have two parameters that need tuning for optimum performance: *C*, which acts as the regularisation parameter, and γ , which decides the degree of influence of a single training example. The models were evaluated, again using the *J*-statistic, across a range of both parameters, producing a grid of evaluation results. Each evaluation was by five-fold cross-validation to penalise parameters that produce over-fitting models. This grid of results is presented as a heatmap and the highest performing parameter combinations are selected for the final models.

RESULTS

Demographics

Ninety four patients with a mean age of 7.3 years were admitted to CUH PICU in this period with a severe TBI. All patients had data

Table 1. Summary of the cohort's demographic and laboratory data, presented for patients split as favourable vs. unfavourable outcomes

Variables	Favourable	Unfavourable	<i>p</i> Value
No. of patients [%]	81 [86]	13 [14]	N/A
Mean age $\pm$ SD (years)	7.3 $\pm$ 5.0	7.9 $\pm$ 5.4	0.67
Males [%]	47 [58]	8 [62]	0.82
Median admission GCS [range]	13 [3–15]	3 [3–10]	0.0000021
Median motor score [range]	6 [1–6]	1 [1–4]	0.0000012
Median Marshall score [range]	1 [1–6]	3 [3–5]	0.022
Mean glucose $\pm$ SD (mmol/L)	7.0 $\pm$ 1.9	10.4 $\pm$ 4.0	0.0045
Mean haemoglobin (g/dL)	10.7 $\pm$ 1.8	11.3 $\pm$ 3.2	0.49
Mean albumin (g/L)	34.7 $\pm$ 4.7	30.0 $\pm$ 7.7	0.051
Mean C-reactive protein (mg/L)	16.9 $\pm$ 23.7	50.2 $\pm$ 43.6	0.010
Mean sodium (mmol/L)	139.7 $\pm$ 3.3	143.7 $\pm$ 7.1	0.024
Mean urea (mmol/L)	4.4 $\pm$ 1.2	4.7 $\pm$ 1.4	0.60
Mean magnesium (mmol/L)	0.79 $\pm$ 0.1	0.8 $\pm$ 0.1	0.81
Mean lactate (mmol/L)	1.5 $\pm$ 0.98	8.2 $\pm$ 4.1	0.00000088
H+ (nmol/L)	39.7 $\pm$ 6.1	57.3 $\pm$ 10.1	0.0000029
Total white cell count ($\times 10^9$ /L)	13.2 $\pm$ 5.7	12.6 $\pm$ 4.6	0.99
Neutrophil count ($\times 10^9$ /L)	10.2 $\pm$ 5.3	9.8 $\pm$ 4.2	0.89
Haematocrit (L/L)	0.31 $\pm$ 0.05	0.29 $\pm$ 0.09	0.36
PT (s)	12.9 $\pm$ 1.4	17.5 $\pm$ 11.3	0.013
APTT (s)	27.8 $\pm$ 3.6	33.1 $\pm$ 9.5	0.030

GCS Glasgow Coma Score, PT prothrombin time, APTT activated partial thromboplastin time, N/A not available

for the 14 variables of interest recorded. The Marshall scores ranged from 1 to 5, with a mode of 2. The cohort's demographic and laboratory data are summarised in Table 1.

Inspection of variables

A graphical plot of all the variables was produced, with the MIC on one axis and the absolute CC on another. Three outstanding variables were identified by visual inspection of the plot. These three variables are lactate, H⁺ (venous), and glucose, and they appear in the top-right quadrant, thereby being the most informative as well as most linearly correlated with the outcome (Fig. 1).

Linear model

Two logistic regression models were trained: a focused model that ignores all blood test variables in the data except the three most

Fig. 1 Maximal information coefficient (MIC) and absolute correlation coefficient (CC) between blood result variables and bipartite outcome. Venous concentrations of lactate, H⁺, and glucose are the most informative variables with outcome

informative ones (i.e. lactate, H⁺ and glucose) and another that considers all the variables. Both models were evaluated with five-fold cross-validation and a decision boundary was chosen to maximise the *J*-statistic.

The focused model (Fig. 2a) predicted favourable outcomes with an AUC = 0.83 (sensitivity = 71%, specificity = 99%). The AUC for the other model (Fig. 2b) in predicting favourable outcomes was 0.90 (sensitivity = 75%, specificity = 99%).

SVM model

Two models were developed: one model considers all the variables and the other considers only the three most informative variables (as previously identified above).

For the three-variable model, the optimal parameters discovered through the grid search (Fig. 3a) produces a model with sensitivity = 80% and specificity = 99%. For the model that considers all variables, optimised parameters through the grid search (Fig. 3b) produces a model with sensitivity = 63% and specificity = 100%.

The performance of the four different models is summarised in Table 2.

DISCUSSION

In this study, we identify admission levels of lactate, H⁺ and glucose as being the most informative of 6-month outcomes after severe paediatric TBI. Using this information, we have constructed two types of model, three-variable and all-variable models, using both logistic regression and SVM analysis to predict outcomes. Of these, the three-variable SVM model was the most accurate in predicting favourable outcomes at 6 months (sensitivity = 80%). However, it is notable that all four models were highly specific (specificity >90% in all cases), that is, reliable in correctly identifying those with unfavourable outcomes.

High admission levels of lactate and glucose, as well as low admission pH, have been previously identified in adult studies as correlating with poor outcomes.^{10–12} Far fewer studies have been conducted in children, but isolated hyperglycaemia has been previously identified in paediatric cohorts as predictive of poor outcomes.^{13–15} However, the therapeutic implications of these findings remain unclear for both adults and children; the most recent international guidelines for managing paediatric TBI found insufficient evidence to recommend intensive glycaemic control and concludes that further research is warranted—findings that are also supported by a recent meta-analysis of ten adult RCTs (*n* = 1066).^{16,17}

Fig. 2 Receiver-operating characteristic (ROC) curves for logistic regression of: **a** The three most informative parameters. Sensitivity = 71%, specificity = 99%. **b** All parameters. Sensitivity = 75%, specificity = 99%

Fig. 3 Heatmaps of the support vector machine (SVM) parameter search. The SVM has two parameters, C (regularisation constant) and γ (decides the degree of influence for any single training example). The performance of the SVM model depends entirely on these parameters. There is no analytical method to deduce the optimal parameters and, as such, they must be chosen by testing a variety of combinations. This variety of combinations is produced on a grid search so that all possible combinations of C and γ , within a set range, are tested. Each point on the heatmap is a unique combination of C and γ . **a** Three most informative parameters. Sensitivity = 80%, specificity = 99%. **b** All parameters. Sensitivity = 63%, specificity = 100%

Table 2. A summary of the four models are presented				
	Sensitivity	Specificity	J-statistic	AUC
Linear model, all variables	0.75	0.99	0.7	0.9
Linear model, three variables	0.71	0.99	0.74	0.83
SVM, all variables	0.63	1	0.79	N/A
SVM, three variables	0.8	0.99	0.63	N/A

SVM support vector machine, AUC area under the curve, N/A not available

Further work is also necessary to elucidate the nature of the relationship between glycaemia and outcomes, as well as to determine whether there is any causal effect underlying the outcome associations and derangements of these three of these metabolic markers. Indeed, caution is advised before drawing any mechanistic conclusions from these correlations; as one example, recent evidence suggests that the injured brain up-regulates lactate metabolism and that a certain degree of extracellular lactate may even be therapeutic after TBI.^{18–20}

The reliability of our models demonstrates the potential for outcome stratification of paediatric TBI patients based on the point-of-admission blood results. This may prove particularly useful given the heterogeneity in outcomes for children compared to adults and would inform both treating clinicians and aid enrolment and classification decisions for clinical trials.²¹ Our model performs favourably compared to existing point-of-admission predictive models in adult TBI patients, with one recent large adult study ($n = 1089$) achieving a sensitivity of 72% in predicting favourable outcomes.^{22–24} Moreover, these models have often included radiological markers from the initial CT scan. A major advantage of our model using only routine laboratory parameters is the accessibility it affords to clinicians in LMICs, where both CT scanners and radiological expertise may be less

readily available. This is particularly important given the global burden of TBI, which is estimated to be three times more commonly encountered in LMIC healthcare systems compared with high-income countries (HIC).²⁵ We envisage that a fully validated model based on laboratory parameters would thus be of use not only in prognostication but also in evaluating the quality of TBI care provision in a healthcare system, as compared to the outcomes achieved in a cohort receiving the best standard of care in a HIC.

More broadly, we have successfully demonstrated a three-part methodology to assess the predictive value of laboratory blood results for a clinical outcome, viz. (i) inspection of variables for the most informative variables, followed by (ii) logistic regression and (iii) SVM analysis. In this way, both linear and complex, interacting, non-linear relationships in a dataset can be discovered. SVM analysis is a state-of-the-art machine-learning algorithm, but it has not yet been widely used in clinical research. However, this is likely to be a developing field, with SVMs being increasingly recognised as a powerful algorithm that is well suited to detecting subtle patterns in ‘noisy’ biological data and a number of proof-of-principle studies have been conducted on large clinical datasets.^{26,27}

Our study was a single-centre review, including only patients admitted to intensive care and as such is limited by a modest sample size. This model must be evaluated on an external database and further work will be required to validate it as a clinical tool; in its current format, it should not be interpreted beyond being a proof-of-principle machine-learning model for serum variables in the context of paediatric TBI. Moreover, in this study we sought to demonstrate the isolated prognostic value of routine serum variables, however, given the statistical significance of admission Glasgow Coma Score and Marshall Score highlighted in our data (Table 1), we would envisage a final, comprehensive clinical tool to also integrate these physiological and imaging parameters. With a larger sample size, it would be valuable also to re-evaluate the model with stratified age ranges, given the known

and suggested differences in pathology after TBI between neonates, infants and older children.⁷ We highlight that early infection and other concurrent injuries are potential confounders that we did not attempt to exclude in this analysis, in order to arrive at a model that is applicable to real-world cohorts. However, this should be taken into account when attempting to draw any mechanistic conclusions about the pathophysiology of TBI based on these results.

CONCLUSION

Prognostication in the context of paediatric TBI is currently very difficult and the development of a reliable predictive model based on point-of-admission data would be of considerable value. We demonstrate a supervised learning method to assess the relationship between 14 admission laboratory variables and outcomes after TBI in children and have developed a proof-of-principle prognostic model using this method.

ACKNOWLEDGEMENTS

Research support—the Medical Research Council (MRC, Grant Nos. G0600986 ID79068 and G1002277 ID98489) and the National Institute for Health Research Biomedical Research Centre (NIHR BRC) Cambridge (Neuroscience Theme; Brain Injury and Repair Theme). A.M.H.Y. is supported by a National Institute for Health Research (NIHR) Academic Clinical Fellowship. P.J.H. is supported by a Research Professorship from the National Institute for Health Research, the NIHR Cambridge Biomedical Research Centre, a European Union Seventh Framework Programme grant (CENTRE-TBI; grant no. 602150), and the Royal College of Surgeons of England.

AUTHOR CONTRIBUTIONS

Conceived and designed the study: S.K., A.M.H.Y., C.M., P.J.H. and S.A. Performed the experiments: S.K., A.M.H.Y. and C.M. Analysed the data: S.K., A.M.H.Y. and C.M. Contributed materials/analysis/data tools: I.J., H.M.F., S.A. and M.R.G. Wrote the paper: S.K. and A.M.H.Y. Critical revision and final approval of paper: I.J., H.M.F., M.R.G., P.J.H. and S.A.

ADDITIONAL INFORMATION

Competing interests: The authors declare no competing interests.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

REFERENCES

- Quayle, K. S. et al. Epidemiology of blunt head trauma in children in U.S. emergency departments. *N. Engl. J. Med.* **371**, 1945–1947 (2014).
- Maas, A. I. R. et al. The Lancet Neurology Commission Traumatic brain injury: integrated approaches to improve prevention, clinical care, and research Executive summary. The Lancet Neurology Commission. *Lancet Neurol.* **16**, 987–1048 (2017).
- Chesnut, R. M. et al. A trial of intracranial-pressure monitoring in traumatic brain injury. *N. Engl. J. Med.* **367**, 2471–2481 (2012).
- Maas, A. I. R., Hukkelhoven, C. W. P. M., Marshall, L. F. & Steyerberg, E. W. Prediction of outcome in traumatic brain injury with computed tomographic characteristics: a comparison between the computed tomographic classification and combinations of computed tomographic predictors. *Neurosurgery* **57**, 1173–1182 (2005). discussion 1173–82.
- Liesemer, K. et al. Use of Rotterdam CT scores for mortality risk stratification in children with traumatic brain injury*. *Pediatr. Crit. Care Med.* **15**, 554–562 (2014).
- Van Beek, J. G. M. et al. Prognostic value of admission laboratory parameters in traumatic brain injury: results from the IMPACT study. *J. Neurotrauma* **24**, 315–328 (2007).
- Figaji, A. A. Anatomical and physiological differences between children and adults relevant to traumatic brain injury and the implications for clinical assessment and care. *Front. Neurol.* **8**, 685 (2017).
- Adelson, P. D. et al. Guidelines for the acute medical management of severe traumatic brain injury in infants, children, and adolescents. Chapter 12. Use of hyperventilation in the acute management of severe pediatric traumatic brain injury. *Pediatr. Crit. Care Med.* **4**, S45–S48 (2003).
- Kochanek, P. M. et al. Guidelines for the acute medical management of severe traumatic brain injury in infants, children, and adolescents—second edition. *Pediatr. Crit. Care Med.* **13**, S1–S2 (2012).
- Jeremitsky, E., Omert, L. A., Dunham, C. M., Wilberger, J. & Rodriguez, A. The impact of hyperglycemia on patients with severe brain injury. *J. Trauma* **58**, 47–50 (2005).
- Griesdale, D. E. G., Tremblay, M.-H., McEwen, J. & Chittock, D. R. Glucose control and mortality in patients with severe traumatic brain injury. *Neurocrit. Care* **11**, 311–316 (2009).
- Stefani, M. A. et al. Elevated glutamate and lactate predict brain death after severe head trauma. *Ann. Clin. Transl. Neurol.* **4**, 392–402 (2017).
- Sayed Saadat, S. M. et al. Association of persistent hyperglycemia with outcome of severe traumatic brain injury in pediatric population. *Childs Nerv. Syst.* **28**, 1773–1777 (2012).
- Kan, C. H., Saffari, M. & Khoo, T. H. Prognostic factors of severe traumatic brain injury outcome in children aged 2–16 years at a major neurosurgical referral centre. *Malays. J. Med. Sci.* **16**, 25–33 (2009).
- Tude Melo, J. R. et al. Mortality in children with severe head trauma: predictive factors and proposal for a new predictive scale. *Neurosurgery* **67**, 1542–1547 (2010).
- Kochanek, P. M. et al. Guidelines for the acute medical management of severe traumatic brain injury in infants, children, and adolescents—second edition author affiliations project management guest editor external peer reviewers. *SI Pediatr. Crit. Care Med.* **13**, S68–S75 (2012).
- Hermanides, J. et al. Glycaemic control targets after traumatic brain injury: a systematic review and meta-analysis. *Crit. Care* **22**, 11 (2018).
- Jalloh, I. et al. Lactate uptake by the injured human brain: evidence from an arteriovenous gradient and cerebral microdialysis study. *J. Neurotrauma* **30**, 2031–2037 (2013).
- Holloway, R. et al. Effect of lactate therapy upon cognitive deficits after traumatic brain injury in the rat. *Acta Neurochir. (Wien.)* **149**, 919–927 (2007).
- Gallagher, C. N. et al. The human brain utilizes lactate via the tricarboxylic acid cycle: a ¹³C-labelled microdialysis and high-resolution nuclear magnetic resonance study. *Brain* **132**, 2839–2849 (2009).
- Lingsma, H. F., Roozenbeek, B., Steyerberg, E. W., Murray, G. D. & Maas, A. I. Early prognosis in traumatic brain injury: from prophecies to predictions. *Lancet Neurol.* **9**, 543–554 (2010).
- Rizoli, S. et al. Early prediction of outcome after severe traumatic brain injury: a simple and practical model. *BMC Emerg. Med.* **16**, 32 (2016).
- Steyerberg, E. W. et al. Predicting outcome after traumatic brain injury: development and international validation of prognostic scores based on admission characteristics. *PLoS Med.* **5**, e165 (2008).
- MRC CRASH Trial Collaborators et al. Predicting outcome after traumatic brain injury: practical prognostic models based on large cohort of international patients. *BMJ* **336**, 425–429 (2008).
- Dewan, M. C. et al. Estimating the global incidence of traumatic brain injury. *J. Neurosurg.* <https://doi.org/10.3171/2017.10.JNS17352>, 1–18 (2018).
- Noble, W. S. What is a support vector machine? *Nat. Biotechnol.* **24**, 1565–1567 (2006).
- Jiang, F. et al. Artificial intelligence in healthcare: past, present and future. *Stroke Vasc. Neurol.* **2**, 230–243 (2017).