

<https://doi.org/10.1038/s44172-025-00352-2>

A mechanics-informed machine learning framework for traumatic brain injury prediction in police and forensic investigations

Check for updates

Yuyang Wei¹, Jeremy Oldroyd², Phoebe Haste³, Jayaratnam Jayamohan⁴, Michael Jones⁵, Nicholas Casey⁶, Jose-Maria Peña⁷, Sonya Baylis⁶, Stan Gilmour^{2,8} ✉ & Antoine Jérusalem¹ ✉

Police forensic investigations are not immune to our society's ubiquitous search for better predictive ability. In the particular and very topical case of Traumatic Brain Injury (TBI), police forensic investigations aim at evaluating whether a given impact or assault scenario led to the clinically observed TBI. This question is traditionally answered by means of forensic biomechanics and neurosurgical expertise which cannot provide a fully objective probabilistic measure. To this end, we propose here a numerical framework-based solution coupling biomechanical simulations of a variety of injurious impacts to machine learning training of police reports provided by the UK's Thames Valley Police and the National Crime Agency's National Injury Database. In this approach, the biomechanical predictions of mechanical metrics such as strain and stress distributions are interpreted by the machine learning model by additionally considering assault specific metadata to predict brain injury outcomes. The framework, only taking as input information typically available in police reports, reaches prediction accuracies exceeding 94% for skull fracture, 79% for loss of consciousness and intracranial haemorrhage, and is able to identify the best predictive features for each targeted injury. Overall, the proposed framework offers new avenues for the prediction, directly from police reports, of any TBI related symptom as required by forensic law enforcement investigations.

Traumatic Brain Injury (TBI) is a pressing public health concern with major social, economic, and medical implications¹⁻³. The incidence of TBI continues to rise, affecting millions of individuals worldwide and resulting in substantial mortality and long-term morbidity⁴⁻⁷. In particular, mild TBI is underreported, challenging to diagnose and linked to long-term neurodegenerative processes⁸⁻¹². As a result, there is an urgent need for accurate assessment tools to predict TBI risk¹³⁻¹⁶. In the particular context of law enforcement forensic investigations, this challenge is further complicated by its judicial implications. Traditionally, this additional dimension is tackled by the involvement of forensic and clinical experts, asked to evaluate whether an injurious scenario may or may not have caused a TBI, a task that is, by definition, not only dependent on the personal assessment of the expert but also on the difficult, if not impossible, quantitative evaluation of said TBI

in probabilistic terms. The development of a reliable and validated simulation environment that can predict the risk of TBI in various assault scenarios is thus of crucial relevance for improving forensic investigations, supporting law enforcement agencies and enhancing public safety¹⁷⁻²¹.

Recent studies have approached this challenge by coupling finite element (FE) models to machine learning. Anderson et al.²² combined FE modelling with network analysis to predict concussion outcomes, highlighting the value of merging biomechanical models with advanced computational methods. Similarly, Cai et al.²³ used deep learning models on brain strain data to classify concussions, showing how machine learning can enhance FE models for predicting TBI severity. Here, we propose a different approach making use of a two-layered machine learning framework to process FE simulation outputs. The work leverages a range of

¹Department of Engineering Science, University of Oxford, Oxford, UK. ²Thames Valley Police, Oxford, UK. ³The Podium Institute for Sports Medicine and Technology, Department of Engineering Science, University of Oxford, Oxford, UK. ⁴John Radcliffe Hospital, Oxford, UK. ⁵School of Engineering, Cardiff University, Cardiff, UK. ⁶National Crime Agency of UK, London, UK. ⁷Lurtis Ltd., Wood Centre for Innovation Stansfeld Park, Oxford, UK. ⁸Keele University, Newcastle, UK. ✉e-mail: stan@oxonadvisory.com; antoine.jerusalem@eng.ox.ac.uk

interdisciplinary expertise in biomechanics, computational mechanics, neurosurgery, neuroimaging and artificial intelligence to predict a set of TBI-related outcomes from a set of inputs, typically provided by a police report for a range of injurious assaults. The resulting framework is structured into two machine learning layers, see Fig. 1. The first layer is a Multilayer Perceptron (MLP) neural network trained against 200 finite element simulations of a range of injurious head impacts. This layer is able to accurately predict a set of maximum strain- and stress-based mechanical quantities in different regions of an idealised head-neck model, without the need to run additional computationally expensive finite element simulations, see Fig. 1b. A second layer, making use of an Extreme Gradient Boosting (XGBoost) algorithm²⁴, is trained with 53 police reports provided by the UK's Thames Valley Police and the UK's National Crime Agency's National Injury Database, see Fig. 1a. Each report was manually post-processed to establish: (i) a kinematic description of the injurious impact and its corresponding boundary conditions, (ii) any relevant additional metadata pertaining to both the victim and the assailant. While both were used as the second layer's inputs, the former was also used by the first layer whose outputs were then additionally fed into the second layer, see Fig. 1c. Following calibration and validation of both layers, the resulting framework achieves high accuracy for key injury types, such as skull fractures, loss of consciousness and intracranial haemorrhages.

The proposed approach also includes a classifier, which allows the identification of the inputs of the framework most relevant for each injury prediction target. Doing so, the framework also identifies which mechanistic quantities, for which head region, has the most injury predictive power. This finding feeds directly into other studies, aimed at identifying the cause and mechanism of head injuries. Finally, it is demonstrated that it is the coupling between these two layers, along with the inclusion of mechanical considerations in the machine learning framework, that enable the high predictive power. By accurately predicting the risk of TBI for different assault scenarios, the research results and findings can help identify high-risk situations, improve risk assessment practices and develop preventive

strategies to mitigate the occurrence and severity of head injuries¹⁹. In the particular case of law enforcement forensic investigation, the proposed simulation-based tool provides law enforcement agencies and forensic medical practitioners with an unprecedented resource for objectively assessing head injuries. Leveraging this tool, investigators can make evidence-based decisions, identify potential suspects and support any subsequent legal process. Furthermore, the findings have important implications for violence prevention efforts, as accurate prediction of TBI risk can inform risk assessment practices and guide the development of targeted preventive strategies.

Results

Biomechanical impact prediction

An FE head-neck model, incorporating a viscoelastic neck support was developed and validated as the mechanical layer of the numerical framework, see Fig. 1b. The FE head model was adapted from a previous version proposed by Schroder et al.²⁵ already validated for head impacts. The viscoelastic neck was modelled with nine pairs of springs and dashpots, connecting the head model to a circular plate. The stiffness of the springs was calibrated against in vivo experimental results from the literature to accurately represent the natural impedance of the human neck. A detailed description of the development and optimisation of the numerical head-neck model can be found in the supplementary material. Note that we purposefully avoid the need for subject-specific models and injury thresholding by relying instead on machine learning to “correct” the generic head model mechanical prediction with metadata²⁶, such as age and gender, to account for individual and situational variations. Taken together, the layer evaluates whether the corresponding mechanical assault leads or not to one of the studied injurious outcomes. Similarly, while impact scenarios involving rigid or elastic objects can be simulated with relative ease, simulating hand-to-head impacts requires a detailed hand model. The model used here was previously developed and validated²⁷ and includes the bone skeleton, subcutaneous tissue and skin, see Fig. S1. To achieve an accurate simulation

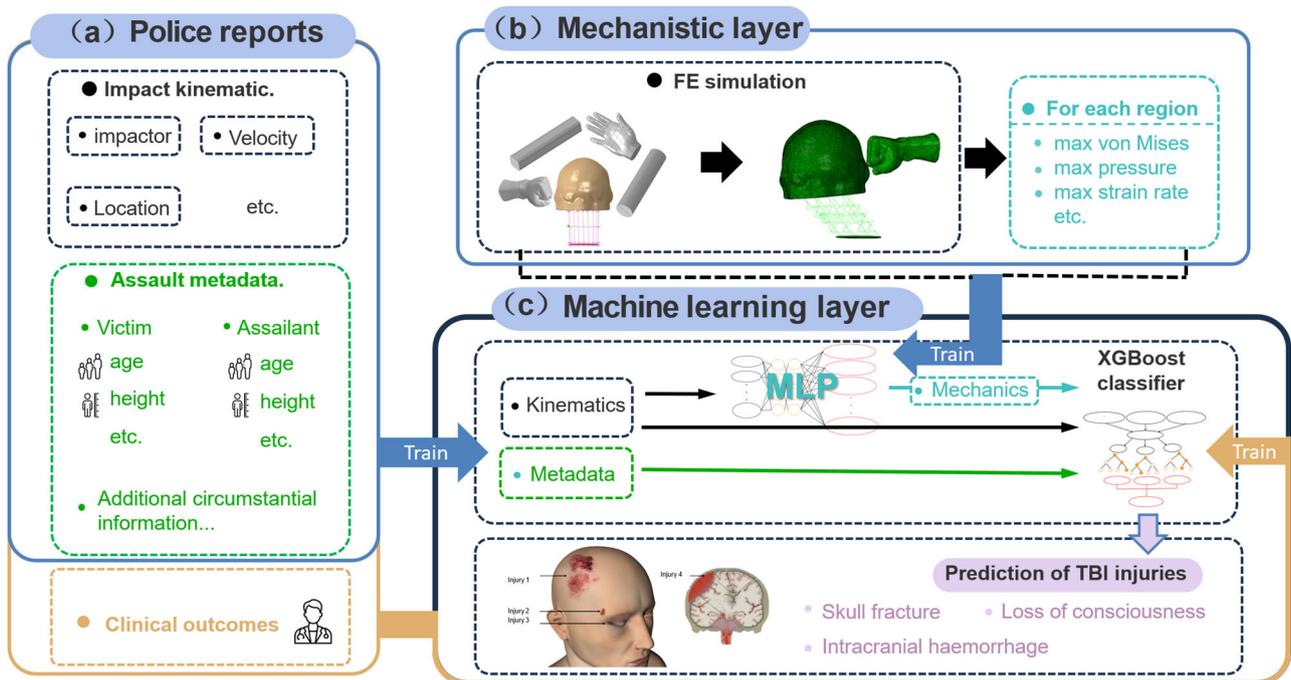


Fig. 1 | Overall framework. a Existing police reports are analysed and three sets of information are extracted: impact kinematic (black), assault metadata (green) and clinical outcomes (ochre); (b) the existing impact kinematic is used to guide a large set of finite element (FE) simulations from which relevant mechanical quantities (blue) are retrieved in different head regions; these inputs and outputs are used to train MLP networks; (c) police report data with clinical outcomes along with MLP

outputs for corresponding impact kinematic information are used to train an XGBoost layer. The final fully trained framework takes as inputs assault metadata, impact kinematic information, and their processing through the MLPs, i.e. the mechanistic “hidden” information, and predicts the risk of TBI injuries (here, the focus is on skull fracture, intracranial haemorrhage and loss of consciousness). A classifier allows the identification of the most relevant inputs for each case.

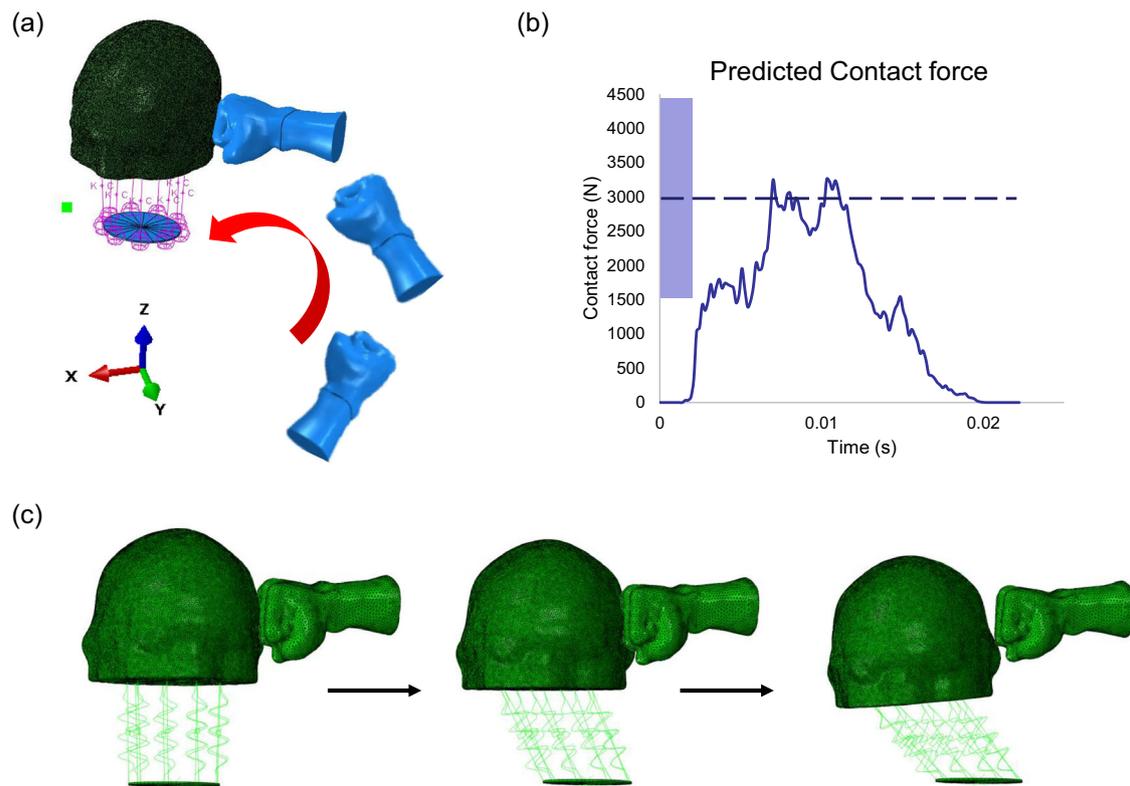


Fig. 2 | Validation of punch simulation using FE head and hand model. a Typical hook punch with the right fist rotating around the left shoulder; (b) predicted contact force during punching is shown, with the range of contact forces measured during

in vivo experiments²⁹ highlighted by the purple bar, and the average value marked by the dashed line. The contact duration can be observed through the variation of contact force over time; (c) simulation before and after impact.

of the specific fist-to-head assaults, the kinematics of the hand was calibrated against previously published experimental measurements conducted on a junior boxer executing a hook punch²⁸, as a good first approximation of what is typically expected from punching occurring in most of our criminal cases. As shown in Fig. 2a, c, a fist velocity of 8.9 m/s, with a predefined deceleration extracted from in vivo experiments²⁹, was applied to the right hand around the rotational centre of the left shoulder. Contact force and duration were used to validate the punching simulation. As shown in Fig. 2b, the measured contact force ranged from 3550 to 5200 N²⁸, and the contact duration from 14 to 28 ms²⁹⁻³¹. The predicted contact force and duration fell within the range of the experimental results.

Along with the punch, other assault conditions including slapping and a range of rigid contact impacts were considered for different impact locations and impact angles, see Fig. 3 and Methods. For each case, impact velocities were varied between 5 and 15 m/s for all impactors. This range was chosen as a first approximation for all impactors and corresponds to the normal punching velocity measured during experiments^{28,30}. Taken altogether, 200 impact simulations were run on the University of Oxford’s Advanced Research Computing service, following a Sobol sequence methodology to sample efficiently the training data of the first machine layer³², see supplementary information on kinematic definitions of assaults, impact scenarios of MLPs and computing platforms and resources for more details. Each simulation provides a set of mechanical outputs for the scalp, skull, grey matter (frontal, occipital, parietal and temporal lobes), thalamus/hypothalamus, white matter as a whole, the white matter of the corpus callosum and the brainstem. The outputs were the maximum von Mises stress, strain rate, pressure, shear energy rate and axonal shear energy rate (only for white matter). These simulations were then used to train the MLP neural networks³³, see Fig. 4a, for each mechanical quantity, i.e. one MLP for the maximum von Mises stress for a subset of the organs, one for the maximum strain rate, etc. Operating as the bridge between raw data and mechanical insights, the MLPs process multifaceted inputs, including the

nature of the assault, impact coordinates, velocity and angle through its hidden layers and output a set of biomechanical quantities to the second machine learning layer to accurately predict different head injuries, see Fig. 4b. The calibration parameters of the MLPs, including learning rates, number of hidden layers and number of neurons on each layer, were optimised to minimise the mean square error between the predicted and original FE simulation results for each mechanical quantity of interest. The full optimisation procedure is detailed in Methods with the final parameters available in Table S1.

FE simulations, while accurate, are computationally intensive and time-consuming, particularly when exploring a wide range of impact scenarios. The MLP neural network, trained on outputs from 200 FE simulations, offers key advantages. First, it substantially reduces computational cost by eliminating the need for rerunning simulations for each new scenario. Once trained, the MLP can quickly predict mechanical quantities, such as strain, stress and their derivatives, for various impact conditions. Second, the MLP can handle a wide range of scenarios, including combinations of impact velocity, angles, and locations that were not part of the training FE simulations data, making it highly scalable.

Injury prediction

The XGBoost algorithm²⁴ is used here to predict head injuries, see Fig. 4b. While it is not limited to these, this second machine learning layer estimates the probability of skull fracture, loss of consciousness and intracranial haemorrhages as the most relevant types of injuries identified in the police reports considered in this work. The algorithm takes as input the kinematic impact description (also used by the MLPs), the outputs of the corresponding MLPs’ predictions and the metadata to predict these specific head injuries. An ensemble of decision trees is constructed, progressively refining their predictive accuracy to evaluate the injury risks associated with each criminal case. The training dataset consists of 53 criminal cases selected from the Thames Valley Police and National Crime Agency’s National

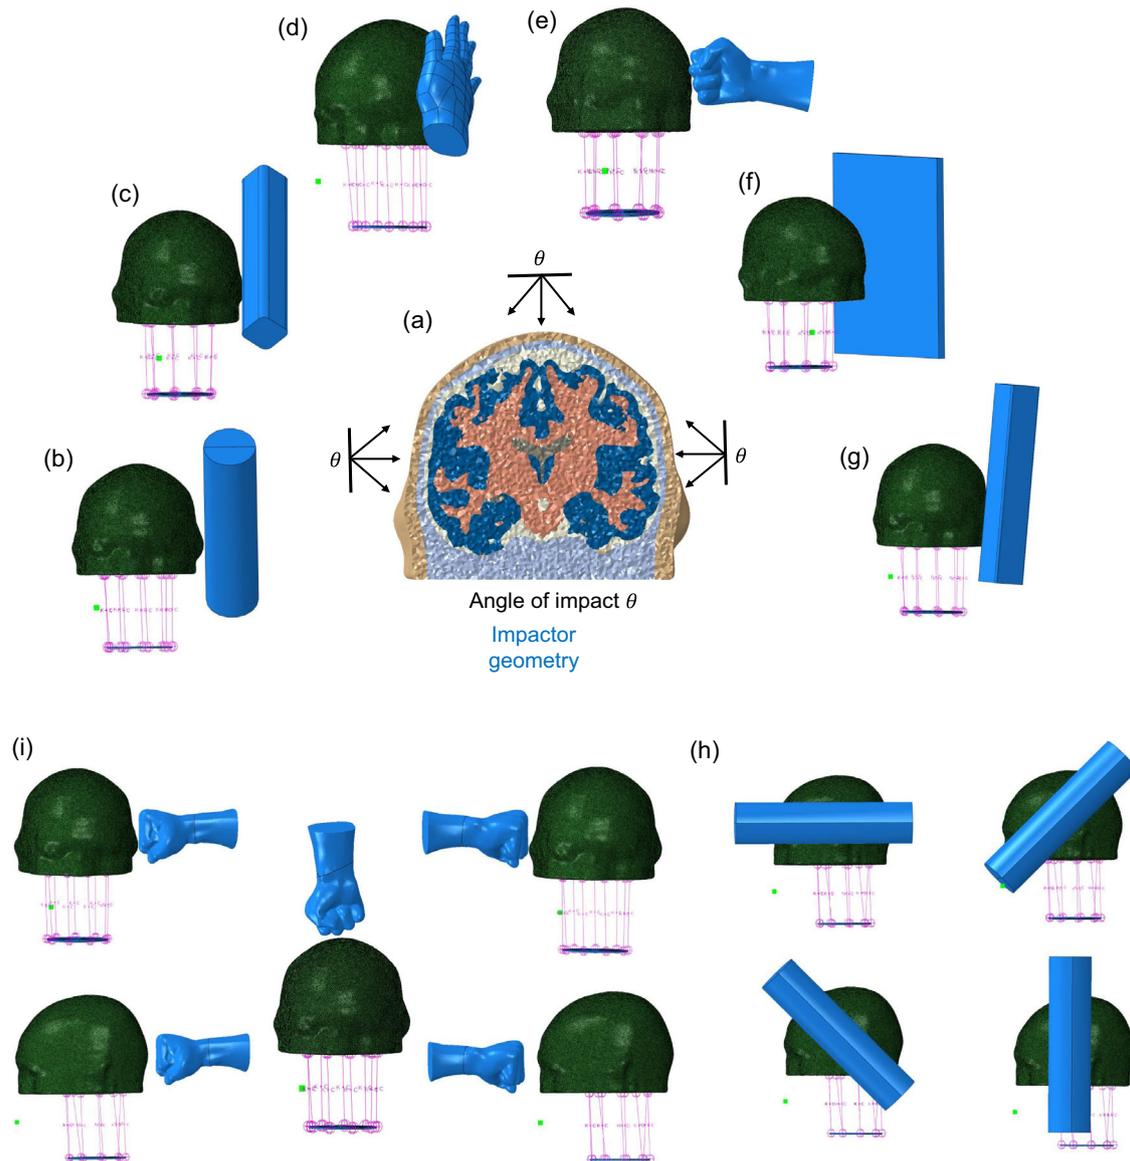


Fig. 3 | Range of impacts considered by the mechanistic layer. a Cross-section view of head model for the three different impact angles considered here; **(b)** cylinder impact; **(c)** blunt corner impact; **(d)** slapping; **(e)** punching; **(f)** flat plate impact; **(g)** sharp corner impact; **(h)** the four different impact orientations considered here; **(i)** the five impact locations considered.

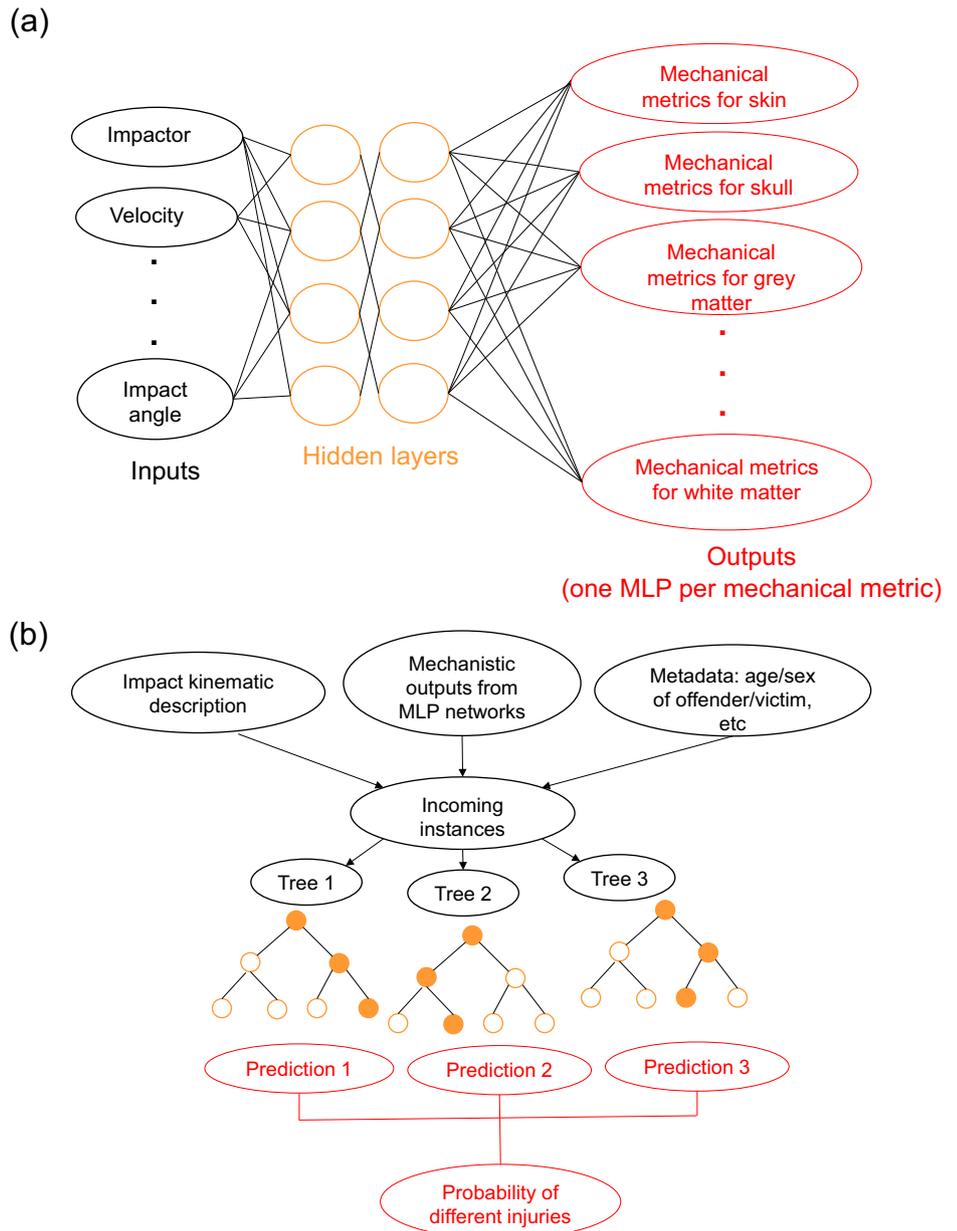
Injury databases. The hyperparameters of XGBoost were optimised using a fivefold algorithm over these 53 cases. To evaluate the accuracy of the second machine learning layer, an additional fivefold cross validation algorithm was used over the 53 cases and the prediction accuracy was calculated across all five folds.

The confusion matrices, combining the test data of all folds, are presented in Fig. 5. The model achieved a very high accuracy of 94% accuracy for skull fracture prediction, a sensitivity of 93% and a specificity of 95%, see Fig. 5a. In total, 14 skull fractures and 36 non-fracture cases were successfully predicted. Feature importance analysis, allowing the identification of which input parameters carry the most predictive value, was also carried out. Skull fracture prediction is provided by averaging the feature importance across all folds. In our framework, feature importance was determined using ensemble tree-based models in Gradient Boosting, which rank input features based on their contribution to predictive performance. The feature importance score is computed by averaging the impurity decrease over all the decision trees within the ensemble³⁴. Features that consistently lead to substantial impurity reduction across multiple trees receive higher importance scores, indicating their strong

influence on model predictions. Figure 5a and Table 2 show that maximum von Mises stress extracted from the scalp played a dominant role in the layer's performance. Other mechanical parameters evaluated by MLPs, as well as some police metadata, also contributed to the predictions, see Discussion below. The simulation framework achieved a prediction accuracy of 79% for loss of consciousness, with a sensitivity of 65% and a specificity of 88%, see Fig. 5b. In other terms, 13 losses of consciousness and 29 non-losses of consciousness cases were successfully predicted when aggregating the folds' tests. The most relevant feature is the maximum pressure in the brainstem, followed by the maximum pressure in the grey matter. Finally, a predicting accuracy of 79% was reached for intracranial haemorrhage, with a sensitivity of 72% and a specificity of 83%, see Fig. 5c. The maximum pressure on brainstem and grey matter are the most important features followed by the equivalent strain on corpus callosum and age of the offender.

Overall, while the accuracy and specificity are quite remarkable, the machine learning feature importance remains unclear without further discussion. It must also be noted that, while both first and second machine learning layers have the capability to deal with a very wide range of inputs,

Fig. 4 | First and second machine learning layers.
a A set of multilayer perceptron (MLP) models, the first machine learning layer (FE surrogate), are used to predict the biomechanical outputs of the FE simulations. The inputs include the type of assault, the location, velocity and angle of the impact, among others. The outputs consist of maximum quantities in different head regions, with each MLP focussing on one mechanical metrics, i.e. von Mises stress, strain rate, pressure and shear energy rate. **b** The eXtreme Gradient Boosting (XGBoost) algorithm, the second machine learning layer, is taking as inputs the kinematic impact description (also used by the MLPs), the outputs of the corresponding MLPs' predictions, and the metadata to predict different head injuries.



the 53 police reports used here do not necessarily (and are unlikely to) encompass all possibilities. In fact, the cases include 33 punches, 1 slap, 3 rigid plate impacts, 5 rigid round impacts, 10 rigid blunt impacts, and 1 rigid sharp impact (see Table S2 for full list). This means that a strong under sampling of many of the inputs is expected, potentially further complicating the analysis of the feature importance at this stage.

Discussion

Framework general performance

The present study integrates advanced FE simulations, machine learning algorithms, and real-case metadata within an interdisciplinary framework to predict the risk of TBI in various assault scenarios within a forensic context. The proposed approach is a major departure from traditional methodologies in criminal investigations³⁵. The results of this study underscore the effectiveness of the simulation platform at predicting TBI-related clinical outcomes with accuracy as high as 94% for skull fracture, 79% for loss of consciousness and traumatic intracranial haemorrhages. In all three cases, both high sensitivities (93%, 65%, and 72%, respectively) and specificities (95%, 88% and 83%, respectively) are observed, albeit

with a markedly lower sensitivity for loss of consciousness. From a law enforcement and forensic medical practitioner perspective, this framework offers new avenues to assess the likelihood of a head injury being sustained under specific circumstances. While this approach does not replace the need for forensic expertise, it offers an objective quantitative contribution to their considerations. Accuracy was assessed using an additional fivefold approach due the relatively small number of cases (53), with all three test datasets used to determine different accuracies, sensitivities and specificities.

Mechanistic contribution to machine learning

One of the key strengths of this framework is the integration of biomechanical simulations within the proposed machine learning algorithms³⁶. This multi-layered approach enhances the accuracy and objectivity of head injury assessments by considering both biomechanical parameters and real-case metadata. Using direct FE simulations alone would not be practical, due to the associated computing costs, including pre-processing, run time and post-processing. Instead, MLPs are employed as surrogates to uncover hidden characteristics not readily available to the overall framework, such as

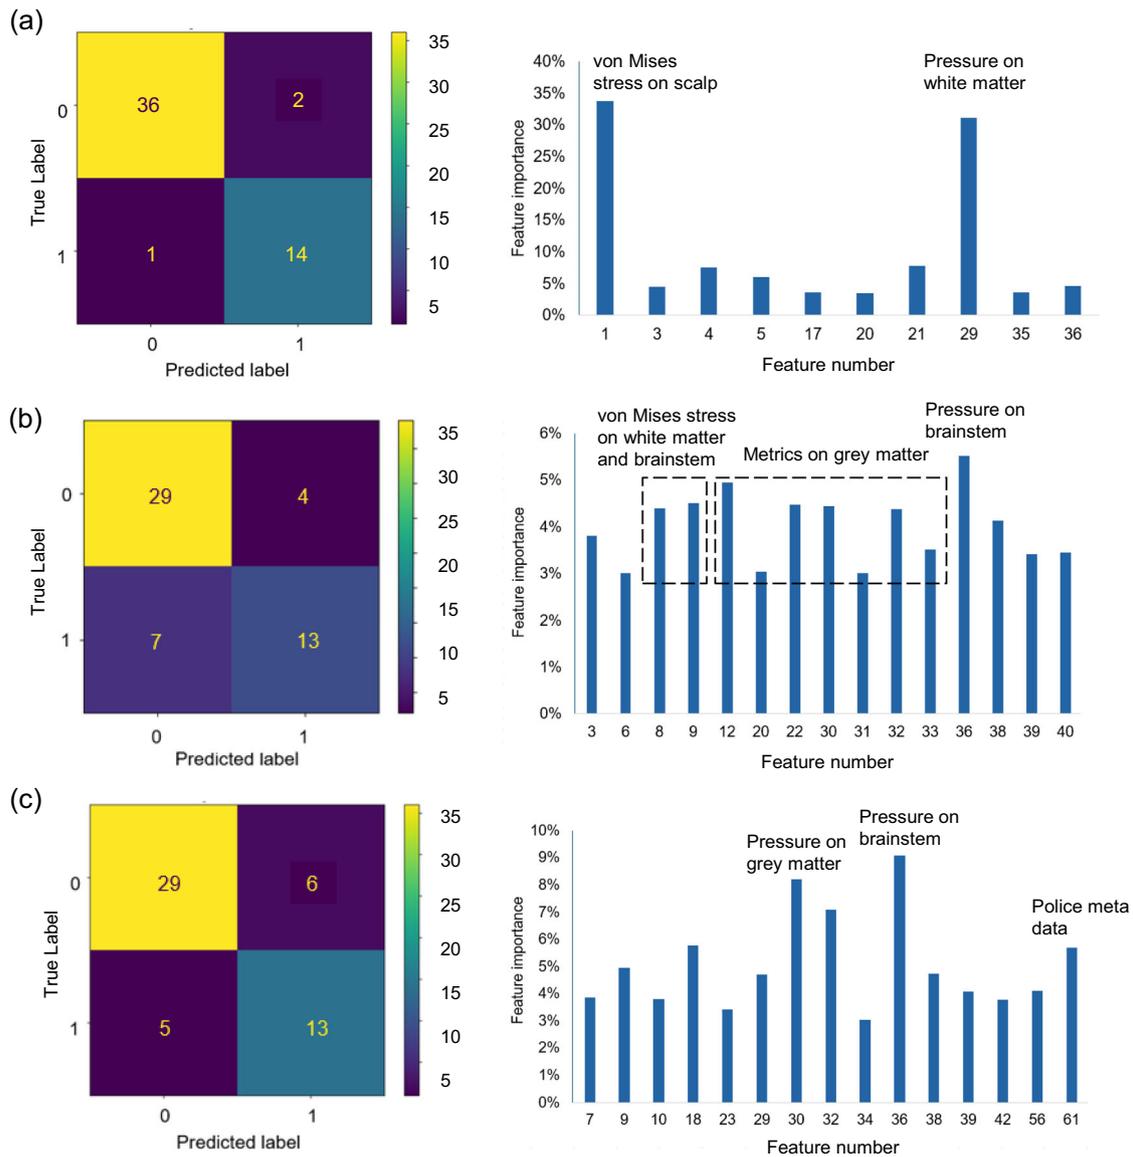


Fig. 5 | Injury prediction results of the numerical pipeline. Confusion matrix and feature importance histogram for (a) skull fractures with 94% accuracy, (b) loss of consciousness with 79% accuracy and (c) for traumatic haemorrhage with 79% accuracy, see Table 2 for feature numbers.

mechanical predictions within the brain. The chosen quantities (maximum von Mises stress, pressure, energy quantities, etc.) span many regions in the head.

Traditionally, criteria based on thresholds at which injuries develop have been proposed^{18,37,38}. However, despite the large body of literature related to the search of a mechanical “silver bullet” threshold for TBI injury, no viable criterion has been, to date, firmly established. This does not imply that additional information related to the mechanical deformation within the brain for a given impact lacks value. In fact, rerunning the XGBoost predictions but without the use of the MLPs (and thus using only the metadata and the impact kinematics as input variables to XGBoost) leads to accuracies of 65%, 52%, 58%, sensitivities of 52%, 49%, 54% and specificities of 77%, 55%, 62% for skull fracture, loss of consciousness and intracranial haemorrhage, respectively, see Table 1. This clearly illustrates that a deeper understanding of mechanics plays a crucial role in the success of the proposed approach. Without this information, accuracy, sensitivity and specificity all suffer, with the sensitivity suffering most. Similarly, not considering metadata in the machine learning also leads to lower metrics, see Table 1. Said otherwise, and with the caveat that other methods could have been used instead of XGBoost, either a pure machine learning layer or

the sole consideration of the FE simulations without metadata are not likely to reach the high accuracy of the full framework proposed here when predicting these injuries. In particular, machine learning alone may lead to a notable number of false negatives, i.e. the lack of mechanical understanding in the predictive layer leads to mistakenly predicting an absence of injury in a few cases.

Machine learning contribution to mechanics

The model achieved an excellent prediction accuracy of 94% for skull fracture, consistently demonstrating high sensitivity and specificity. This remarkable performance can be attributable to the strong correlation between mechanical quantities, derived from biomechanical simulations and the likelihood of a skull fracture occurring³⁹. Paradoxically, the simulation framework identifies the scalp’s von Mises stress as the best predictive feature for this, while one would have expected the skull to be selected. To understand this apparent contradiction, the Pearson correlation coefficients⁴⁰ are calculated here for the 200 simulation FE inputs and outputs generated for the training of the MLPs. These coefficients provide a statistical measure, quantifying linear relationships between two variables: a value close to 1 implies a strong positive linear correlation, while a value of 0

Table 1 | Comparisons of accuracy, area under the curve (AUC), sensitivity and specificity with and without Multilayer Perceptron (MLP) neural networks or metadata for prediction of skull fracture, loss of consciousness and intracranial haemorrhage

Injury	Accuracy	AUC	Sensitivity	Specificity
Skull fracture with MLP	94%	94%	93%	95%
Skull fracture without MLP	65%	68%	52%	77%
Skull fracture without metadata	83%	86%	71%	89%
Loss of consciousness with MLP	79%	78%	65%	88%
Loss of consciousness without MLP	52%	58%	49%	55%
Loss of consciousness without metadata	66%	67%	50%	76%
Intracranial haemorrhage with MLP	79%	80%	72%	83%
Intracranial haemorrhage without MLP	58%	70%	54%	62%
Intracranial haemorrhage without metadata	68%	76%	63%	71%

Table 2 | List of inputs for second machine learning layer (XGBoost); mechanical inputs (outputs of the Multilayer Perceptron (MLP) neural networks: maximum value), impact kinematic, police metadata

Feature number	Inputs	Feature number	Inputs
1	Scalp, von Mises Stress	32	Grey matter (parietal lobe), pressure
2	Skull, von Mises Stress	33	Grey matter (temporal lobe), pressure
3	Grey matter (frontal lobe), von Mises Stress	34	Thalamus/hypothalamus, pressure
4	Grey matter (occipital lobe), von Mises Stress	35	White matter, pressure
5	Grey matter (parietal lobe), von Mises Stress	36	Brainstem, pressure
6	Grey matter (temporal lobe), von Mises Stress	37	Corpus callosum, pressure
7	Thalamus/hypothalamus, von Mises Stress	38	Thalamus/hypothalamus, axonal shear energy rate
8	White matter, von Mises Stress	39	White matter, axonal shear energy rate
9	Brainstem, von Mises Stress	40	Brainstem, axonal shear energy rate
10	Corpus callosum, von Mises Stress	41	Corpus callosum, axonal shear energy rate
11	Grey matter (frontal lobe), equivalent strain	42	Punching
12	Grey matter (occipital lobe), equivalent strain	43	Slapping
13	Grey matter (parietal lobe) equivalent strain	44	Impact from a plate
14	Grey matter (temporal lobe), equivalent strain	45	Impact from a round object
15	Thalamus/hypothalamus, equivalent strain	46	Impact from a blunt object
16	White matter, equivalent strain	47	Impact from a sharp object
17	Brainstem, equivalent strain	48	Impact location (factor 1)
18	Corpus Callosum, equivalent strain	49	Impact location (factor 2)
19	Grey matter (frontal lobe), shear energy rate	50	Impact angles
20	Grey matter (occipital lobe), shear energy rate	51	Impact orientations
21	Grey matter (parietal lobe), shear energy rate	52	Impact velocity
22	Grey matter (temporal lobe), shear energy rate	53	<u>Gender (Victim)</u>
23	Thalamus/Hypothalamus, shear energy rate	54	<u>Body built (Victim)</u>
24	White matter, shear energy rate	55	<u>Height (Victim)</u>
25	Brainstem, shear energy rate	56	<u>Age (Victim)</u>
26	Corpus callosum, shear energy rate	57	<u>Previous TBI (Victim)</u>
27	Scalp, pressure	58	<u>Gender (Offender)</u>
28	Skull, pressure	59	<u>Body built (Offender)</u>
29	Corpus Callosum, pressure	60	<u>Height (Offender)</u>
30	Grey matter (frontal lobe), pressure	61	<u>Age (Offender)</u>
31	Grey matter (occipital lobe), pressure		

implies no correlation. The list of variables considered here, including kinematics data and all the corresponding mechanical data outputted by the MLPs, i.e. the inputs to XGBoost, is provided in Table 2, and the resulting correlation heatmap is shown in Fig. 6a. One notable observation is

that all mechanical quantities (features <42) exhibit relatively strong correlations with each other and with the impact velocity (feature 52). Regarding skull fracture specifically, von Mises stress of both the scalp (feature 1) and skull (feature 2) are highly correlated for each respective

quantity, and they are strongly correlated across different regions of the head. This indicates that while the proposed layer selected the scalp, it equally correlates with the skull. This can be rationalised as follows: scalp stress reflects the mechanical forces absorbed and distributed by the scalp, in turn transferred to the underlying skull. This correlation suggests that scalp stress can thus serve as an early indicator of forces likely to cause skull fractures. In the training of the XGBoost model, feature importances are derived from the ensemble of decision trees used to make predictions. While scalp von Mises stress was ranked higher in importance than skull von Mises stress across a range of impact scenarios, the close proximity of both feature and the consistent correlation of stress in both regions for the loading conditions considered here rationalise the inability of the tree to differentiate between them. As such, while XGBoost is particularly effective at handling correlated variables by selecting the most relevant features for the prediction task while accounting for their interdependence, correlation study between features remains a required step to infer causation from correlations. This finding aligns with biomechanical observations, indicating that bone fractures will occur when the stress exceeds a given threshold. That threshold can go as low as 25 MPa in children⁴¹ and has been previously reported as high as 92.71 MPa⁴². The advantage of the proposed framework is that, while the proposed FE simulations and resulting MLP layers are done for a generic head not necessarily representative of the victim of interest, the predicted values of maximum von Mises stress for the skull can be incorporated in XGBoost and their meaning (in terms of risk of fracture) can then be “decided” by XGBoost by accounting for additional metadata information, such as the age of the victim. Obviously, this additional distinction requires a larger training dataset if one wants to account for all possible age groups (we only had adult victims in our dataset).

The model accuracy for predicting loss of consciousness reaches 79%. According to the feature importance analysis, the maximum pressure in all four lobes (features 30–33 in Table 2) is strongly correlated with each other and with the pressure in the brainstem (feature 36 in Table 2), which emerged as the most predictive feature. This was followed by the pressures in all the grey matter lobes. The strongest candidates for predicting loss of consciousness are the mechanical parameters in the brainstem, which is consistent with medical findings that attribute loss of consciousness primarily to brainstem dysfunction. This aligns with the understanding that the brainstem, playing a crucial role in maintaining consciousness by regulating wakefulness and attention, as well as maintaining homeostasis, is particularly vulnerable to increased pressure during traumatic impacts. Elevated brainstem pressure can disrupt the reticular activating system, resulting in temporary or prolonged loss of consciousness depending on the severity and location of the impact^{43–46}. While the accuracy for predicting loss of consciousness is still considerable, it is notably lower than that achieved for skull fractures. This disparity highlights the intricate nature associated with this injury, which can be influenced by various factors, such as the location and angle of impact, as well as individual variations in response to trauma, including factors such as the neck’s width of the victim or the victim’s state, e.g. during alcohol intoxication^{26,47,48}. Based on our analysis of criminal reports from Thames Valley Police and the National Crime Agency, it was found that cases of loss of consciousness were more likely to occur when the victims were under the influence of alcohol or drugs. To further enhance the accuracy of predicting this outcome, future research may consider refining variables, such as the victim’s health status and prior medical conditions. For this research, it was assumed that all victims had no previous TBI unless clarified in the criminal reports. While the mechanistic rationale remains arguably unclear, it is, however, remarkable that the brainstem (and immediate neighbouring regions) mechanical metrics are consistently involved in the prediction, either directly or through correlative effects, to loss of consciousness, in agreement with the current clinical understanding of the role of brainstem in such outcome⁴⁹. XGBoost naturally handles multicollinearity through its decision tree structure, which selects features based on impurity reduction. However, multicollinearity can still influence feature importance rankings, as correlated features may share predictive power. To address this, the Pearson correlation analysis (Fig. 6a) grouped

highly correlated features, such as brainstem pressure (Feature 36) with nearby mechanical features contributing to similar outcomes. XGBoost’s L1 and L2 regularisation helped manage model complexity and prevent overfitting, ensuring that no single feature dominated the rankings. While multicollinearity affects individual rankings, the collective contribution of related features is reflected, with brainstem pressure emerging as a key predictor for loss of consciousness.

The proposed machine learning framework demonstrates substantial promises in predicting various head injury outcomes resulting from assaults⁵⁰. While impressive accuracy was achieved in certain scenarios, the complexity of head injuries and the influence of individual variations requires continued research and model refinement to fully unlock the potential of this approach in aiding criminal investigations. These findings prompt several conjectures. Firstly, the accuracy achieved in predicting skull fractures suggests that the force and velocity parameters employed in the biomechanical simulations play a pivotal role in determining the severity of head injuries. This correlation between impact severity and injury outcome is a key takeaway from the study. However, the relatively lower accuracy in predicting loss of consciousness underscores the complexity of this outcome, though the model demonstrated its ability at linking it to mechanical stress at and around the brainstem. Lastly, the accuracy achieved in predicting intracranial haemorrhages highlights the model’s sensitivity to factors related to bleeding within the brain. This research stands as a remarkable collaboration between academia, law enforcement and medical experts. Through dynamic partnerships with forensic specialists and investigators, we have ensured the tool’s real-world viability and transformative potential. The collaboration with law enforcement agencies, such as the Thames Valley Violence Reduction Unit and the National Crime Agency-National Injury Database, is instrumental in validating the simulation tool with real case scenarios. This integration of real-world data adds practical relevance to the findings and ensures the tool’s applicability in criminal investigations, potentially aiding in the identification of suspects and supporting the prosecution of perpetrators.

In this study, direct mechanical inputs from FE simulations, such as pressure and axonal shear energy rate, were used instead of composite postprocessed metrics like the Brain Injury Criterion or the Head Injury Criterion⁵¹. Doing so, we aim at using direct FE outputs and let the machine learning layer interpret their significances based on additional metadata. Using direct mechanical quantities allows the model to better capture the brain’s response to impacts in real-time. This approach avoids the limitations of predefined thresholds and enables the machine learning model to autonomously learn from the data, improving its ability to generalise and predict traumatic brain injuries more effectively.

Linking mechanics and injury

The complexity related to the many mechanical quantities and their correlations can be alleviated—at least from an analysis perspective—by using the DALEX library⁵². This library provides a framework for understanding variable importance in machine learning models. It calculates variable single or group importance by systematically perturbing feature values, while keeping all other factors constant, and observing the resultant change in model performance. By permuting these grouped features and measuring the corresponding drop in predictive accuracy, their collective importance can be quantified. Features with higher scores indicate a more substantial impact on the model’s predictions, while features with lower scores contribute less to the overall predictive power of the model.

By categorising variable importance into groups based on Pearson correlation coefficients and feature importance analysis, we identified the potential candidates for predicting the three injuries according to the results shown above: (i) pressure on the brainstem, (ii) pressure on the grey matter regions, (iii) von Mises stress on the skull, and (iv) metadata (ages and genders of offenders and victims). As depicted in Fig. 6b–d, insights into the model’s dynamics were gained by examining grouped variable importance, where higher scores indicate a greater influence on the model’s predictions.

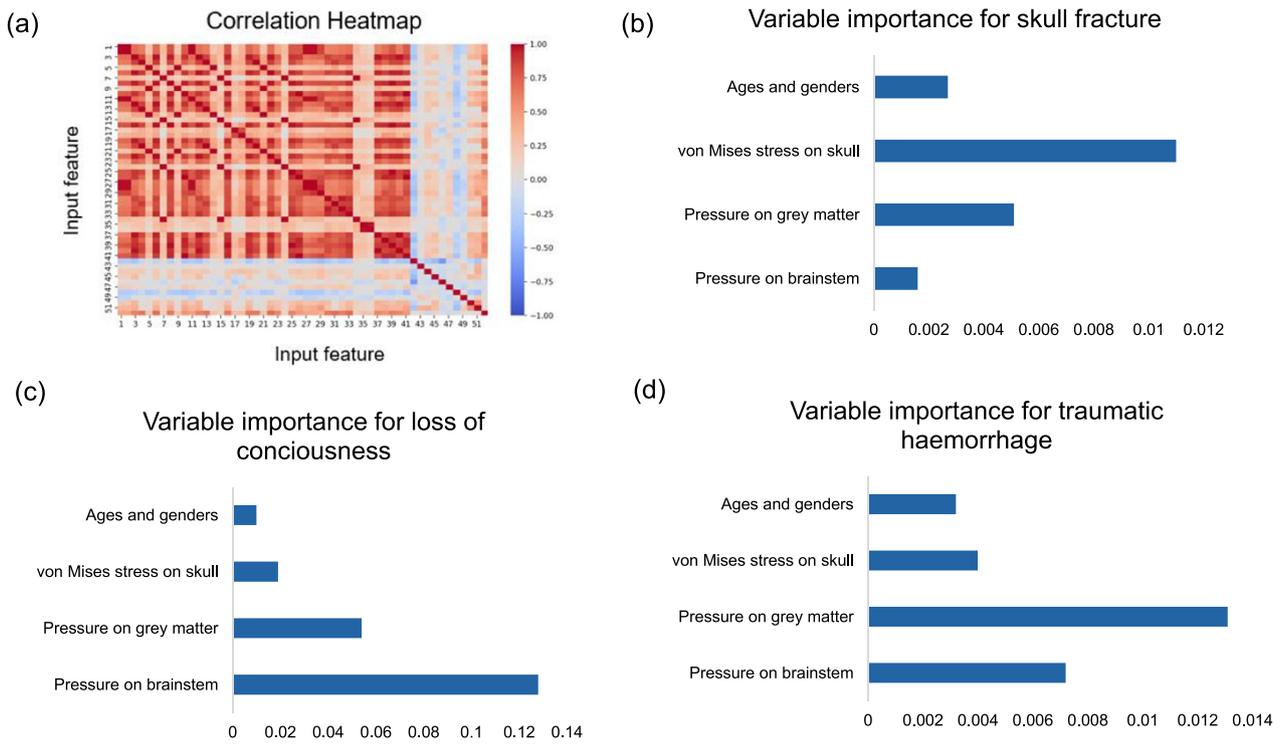


Fig. 6 | Feature analysis. a Pearson coefficients among all pairs of non-metadata of inputs for eXtreme Gradient Boosting (XGBoost), see Table 2 for feature numbers; variable importance for skull fracture (b), loss of consciousness (c) and traumatic intracranial haemorrhage (d). The x-axes in (b), (c), and (d) represents the variable

importance calculated with the DALEX package by finding the top N data points closest to the explained instance (variable drop-out), performing permutation variable importance and using 1-AUC as loss function.

Specifically, the mechanical values associated with the skull play a dominant role in achieving high prediction accuracy for skull fractures. Conversely, pressure on the brainstem is the most relevant parameter for predicting loss of consciousness. For predicting traumatic haemorrhage, pressure on the grey matter is the most important factor. This analysis reaffirms the earlier findings, emphasising the strong dependence of skull fractures on mechanical metrics related to the skull, while loss of consciousness is predominantly influenced by metrics associated with the brainstem. Moreover, traumatic haemorrhages are affected by metrics from both grey matter and surrounding regions, aligning with the expected injury locations, e.g. bridging veins.

Methods

Impact FE simulations

Each FE simulation is run as a dynamic explicit simulation on Abaqus 2022 software (see supplementary information on computing platforms and resources), allowing for the tracking of the time evolution of various mechanical metrics within the different regions of the head. Simulations were defined by different inputs, as are shown in Table 2, including the different types of assaults, impact velocity, angle, location of incidence and the impactor geometry. Except for the hand-to-head impacts, all impactors within the simulations were represented as rigid bodies and a friction coefficient of 0.4 was applied, consistent with previous research²⁵. Four distinct rigid impactor geometries were employed to encompass a wide range of possible impact scenarios. The first, referred to as the round impactor, assumed a cylindrical shape with a radius of curvature of 3.6 cm. The second, a blunt impactor, featured a right-angled analytical surface smoothed along the edge, incorporating a quarter of a cylinder with a 1 cm radius of curvature; the sharp impactor was modelled as a right angle smoothed with an edge of 0.3 cm; lastly, a flat plate was also considered. Additional information is provided as supplementary information on the kinematic definitions of assaults.

Mechanistic machine learning layer: MLPs

MLP networks are used here to predict specific mechanistic quantities (as a set of maximum values for each region) calculated otherwise by means of FE simulations at a much larger computational cost. Each takes as inputs specific boundary conditions pertaining to an impact of interest (see Table S3) and outputs the set of regional quantities for the metric it is responsible for (e.g. von Mises stress, pressure, etc.). Each MLP is made of multiple neuron layers; the first one of which consists of an array of input parameters characterising the assault scenario, such as assault type, location, velocity and impact angle. The MLPs can then have multiple hidden layers, each composed of interconnected neurons designed to learn and transform input information effectively. These hidden layers introduce nonlinearity through activation functions⁵³. The output layer finally provides a final prediction of the injury that the MLP is being trained for. Following the general form of the forward pass of the MLP, each j th neuron of the $i+1$ th layer receives a value Y_{j+1}^i calculated from the values carried by the previous layer's neurons.

$$Y_{j+1}^i = f(W_i \cdot X_i + b_{i+1}^j) \quad (1)$$

where the input X_i is the vector of values carried by the previous layer's neurons, W_i is the vector of corresponding weights and whose product is passed through the activation function f after the incorporation of a bias b_{i+1}^j . The MLP's learning process adjusts all weights and biases during training to minimise the prediction error and improve the accuracy of TBI risk assessment. It is done here by backpropagation, implemented through Stochastic Gradient Descent⁵⁴.

The MLP model was chosen for its computational advantages over traditional FE simulations. Specifically, the MLP substantially reduces computational cost and time by predicting mechanical quantities such as strain and stress, without the need for rerunning extensive FE simulations for each new scenario. Although a trade-off between computational

efficiency and prediction accuracy can be predicted, the MLP achieved c. 90% of prediction accuracy with an optimised set of hyperparameters. This was achieved with a grid search technique over fivefold. This exploration allows for tuning of the learning rate (0.001, 0.01 and 0.1), determining the number of hidden layers (ranging from 1 to 5) and configuring the number of neurons on each layer (ranging from 2 to 36). The process seeks to minimise the mean square error between predicted and actual values. Table S1 shows the final combinations of hyperparameters for each MLP along with the errors associated to them.

Note finally that various algorithms were tested to assess their effectiveness in predicting mechanical properties from the FE simulation data. The tested algorithms included MLP, Bayesian Neural Networks (BNN), Deep Neural Networks (DNN), Support Vector Machines (SVM), and k-Nearest Neighbours (KNN). Among these, MLP demonstrated the best performance, achieving the lowest relative error between the predicted and real data, making it the optimal choice for this layer, see Table S4. The optimised hyperparameters for all these different algorithms are presented in Tables S5–8. Additional information is provided in the supplementary information on impact scenarios of MLPs.

Injury machine learning layer: XGBoost

The second machine learning layer in the framework leverages XGBoost, a gradient-boosting algorithm known for its efficiency and predictive performance²⁴. In this layer, the inputs include the mechanical predictions, such as maximum regional stresses and powers predicted by the MLPs along with the MLPs' inputs that led to those, as well as the metadata collected from the 53 police reports, containing information such as the sex, age, height and weight of the offender or victim, among other factors. These reports were postprocessed for use in the pipeline, including identification, data cleaning, outlier detection, handling of missing values, anonymisation. XGBoost sequentially constructs an ensemble of decision trees, with each tree trained to correct the errors of its predecessor. It then aims to minimise a loss function, typically mean squared error for regression tasks or log loss for classification by iteratively adding decision trees²⁴. The final prediction is obtained as the weighted sum of these individual tree predictions:

$$y_i = \phi(x_i) = \sum_{j=1}^K f_j(x_i) \quad (2)$$

Here, y_i represents the predicted TBI injury for individual x_i , K is the total number of trees in the ensemble and f_j denotes the prediction contribution of the j th tree. By making use of a grid search approach, a range of hyperparameter configurations to identify the optimal settings for our specific biomedical dataset was systematically explored, see final set of hyperparameters in Table S9. To evaluate the model's performance, a k -fold cross-validation strategy was employed. This means dividing the 53 cases in $k = 5$ subsets, training the XGBoost model on $k-1$ folds and validating it on the remaining fold. This process was repeated k times, with each fold serving as the validation set exactly once. The cross-validation results provided a robust estimate of the model's generalisation performance and aided at identifying potential issues such as overfitting.

In assessing the predictive capabilities of the present XGBoost model, a battery of performance metrics was considered. These included classification accuracy, precision, recall, F1-score and area under the receiver operating characteristic curve. The choice of these metrics was deliberate, as they collectively offered a comprehensive view of the model's strengths and limitations, especially concerning its ability to handle class imbalances—a frequent challenge in biomedical datasets. Note that other algorithms were also evaluated, including BNN, DNN, Random Forest, XGBoost, SVM, KNN, and Logistic Regression. XGBoost was selected as the best-performing model, as it achieved the highest prediction accuracy across the three injuries being studied, see Tables S10–12. The optimised hyperparameters for all these different algorithms are presented in Tables S13–16.

The dataset consists of 53 criminal cases, which is not a large number due to the difficulty in obtaining reliable witness recordings and medical assessments. Consequently, an independent dataset was not used to objectively evaluate the machine learning model. However, K -fold cross-validation with $K = 5$ was employed to rigorously test the model on data not used during training, enhancing its robustness by preventing contamination between training and validation sets. As such, each subset of data serves as validation at least once, ensuring the model is evaluated on diverse portions. This approach provides a fair and reliable measure of performance and substantially reduces overfitting, compensating for the lack of a completely independent dataset.

Processing of the police reports

The dataset was compiled from detailed police reports, which were post-processed and curated by a police officer to ensure accuracy and consistency. All reports and data were anonymised by Thames Valley Police and the National Crime Agency before aggregation for analysis. Each case was described based on multiple sources, including CCTV footage and witness testimonies. The impact severity was classified into three levels: low, medium, and heavy, determined by the evidence provided. Additionally, the impact speed was estimated using the offender's height, weight, and the assigned impact level. Further details on these criteria can be found in the supplementary information. To protect privacy, all the personal information, such as names or any identifiable characteristic, was anonymised, while key demographic data (age and gender) were retained for analysis. Any inconsistencies in the reports, such as conflicting witness accounts or unclear footage, were addressed through further review by the police officers to reduce noise in the dataset and ensure data accuracy. This approach aligns with recent forensic methodologies, such as the use of subject-specific FE head models for skull fracture evaluation in forensic pathology, as discussed by Henningsen et al.⁵⁵. Inclusion and exclusion criteria were applied to select cases with sufficient detail for analysis, with incomplete cases excluded.

Detailed information about specific injuries and related impact scenarios was not provided to further avoid attempts at identifying each case. However, a detailed breakdown of impact types, angles, and the distribution of male and female victims and offenders is available in the dataset (Tables S2 and S3).

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

The datasets generated during and/or analysed during the current study are not publicly available, to maximise patients/victims' confidentiality; instead, a dummy dataset is provided for illustration.

Code availability

The Python codes for the machine learning layers, along with the related dummy metadata, are available on OSF with the <https://doi.org/10.17605/OSF.IO/6XB2P>.

Received: 23 June 2024; Accepted: 21 January 2025;

Published online: 26 February 2025

References

- Hyder, A. A., Wunderlich, C. A., Puvanachandra, P., Gururaj, G. & Kobusingye, O. C. The impact of traumatic brain injuries: a global perspective. *NeuroRehabilitation* **22**, 341–353 (2007).
- Maas, A. I., Stocchetti, N. & Bullock, R. Moderate and severe traumatic brain injury in adults. *Lancet Neurol.* **7**, 728–741 (2008).
- Corrigan, J. D., Selassie, A. W. & Orman, J. A. L. The epidemiology of traumatic brain injury. *J. Head. Trauma Rehabil.* **25**, 72–80 (2010).
- Dewan, M. C. et al. Estimating the global incidence of traumatic brain injury. *J. Neurosurg.* **130**, 1080–1097 (2018).

5. Feigin, V. L. et al. Incidence of traumatic brain injury in New Zealand: a population-based study. *Lancet Neurol.* **12**, 53–64 (2013).
6. Marin, J. R., Weaver, M. D., Yealy, D. M. & Mannix, R. C. Trends in visits for traumatic brain injury to emergency departments in the United States. *JAMA* **311**, 1917–1919 (2014).
7. Peeters, W. et al. Epidemiology of traumatic brain injury in Europe. *Acta Neurochir.* **157**, 1683–1696 (2015).
8. Giza, C. C. et al. Summary of evidence-based guideline update: evaluation and management of concussion in sports: report of the Guideline Development Subcommittee of the American Academy of Neurology. *Neurology* **80**, 2250–2257 (2013).
9. Cassidy, J. D. et al. Incidence, risk factors and prevention of mild traumatic brain injury: results of the WHO Collaborating Centre Task Force on Mild Traumatic Brain Injury. *J. Rehabil. Med.* **36**, 28–60 (2004).
10. Gardner, R. C. & Yaffe, K. Epidemiology of mild traumatic brain injury and neurodegenerative disease. *Mol. Cell. Neurosci.* **66**, 75–80 (2015).
11. Ruff, R. M. et al. Recommendations for diagnosing a mild traumatic brain injury: a National Academy of Neuropsychology education paper. *Arch. Clin. Neuropsychol.* **24**, 3–10 (2009).
12. McKee, A. C. et al. Chronic traumatic encephalopathy in athletes: progressive tauopathy after repetitive head injury. *J. Neuropathol. Exp. Neurol.* **68**, 709–735 (2009).
13. Duhaime, A.-C. et al. Common data elements in radiologic imaging of traumatic brain injury. *Arch. Phys. Med. Rehabil.* **91**, 1661–1666 (2010).
14. Shenton, M. E. et al. A review of magnetic resonance imaging and diffusion tensor imaging findings in mild traumatic brain injury. *Brain Imaging Behav.* **6**, 137–192 (2012).
15. Bazarian, J. J., Blyth, B., Mookerjee, S., He, H. & McDermott, M. P. Sex differences in outcome after mild traumatic brain injury. *J. Neurotrauma* **27**, 527–539 (2010).
16. Narayan, R. K. et al. Clinical trials in head injury. *J. Neurotrauma* **19**, 503–557 (2002).
17. Post, A., Hoshizaki, B. & Gilchrist, M. D. Finite element analysis of the effect of loading curve shape on brain injury predictors. *J. Biomech.* **45**, 679–683 (2012).
18. Zhang, L., Yang, K. H. & King, A. I. A proposed injury threshold for mild traumatic brain injury. *J. Biomech. Eng.* **126**, 226–236 (2004).
19. Ji, S. et al. Parametric comparisons of intracranial mechanical responses from three validated finite element models of the human head. *Ann. Biomed. Eng.* **42**, 11–24 (2014).
20. Mao, H. et al. Development of a finite element human head model partially validated with thirty five experimental cases. *J. Biomech. Eng.* **135**, 1111002 (2013).
21. Lindgren, N. et al. Prediction of skull fractures in blunt force head traumas using finite element head models. *Biomech. Model. Mechanobiol.* **23**, 207–225 (2024).
22. Anderson, E. D., Giudice, J. S., Wu, T., Panzer, M. B. & Meaney, D. F. Predicting concussion outcome by integrating finite element modeling and network analysis. *Front. Bioeng. Biotechnol.* **8**, 309 (2020).
23. Cai, Y. et al. Concussion classification via deep learning using whole-brain white matter fiber strains. *PLoS ONE* **13**, e0197992 (2018).
24. Chen, T. & Guestrin, C. *Proc. 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.* 785–794.
25. Schroder, A. et al. A machine learning enhanced mechanistic simulation framework for functional deficit prediction in TBI. *Front. Bioeng. Biotechnol.* **9**, 587082 (2021).
26. Kleiven, S. & Von Holst, H. Consequences of head size following trauma to the human head. *J. Biomech.* **35**, 153–160 (2002).
27. Wei, Y., Zou, Z., Wei, G., Ren, L. & Qian, Z. Subject-specific finite element modelling of the human hand complex: muscle-driven simulations and experimental validation. *Ann. Biomed. Eng.* **48**, 1181–1195 (2019).
28. Dinu, D. & Louis, J. Biomechanical analysis of the cross, hook, and uppercut in junior vs. elite boxers: implications for training and talent identification. *Front. Sports Act. Living* **2**, 598861 (2020).
29. Atha, J., Yeadon, M., Sandover, J. & Parsons, K. The damaging punch. *Br. Med J.* **291**, 1756–1757 (1985).
30. Walilko, T. J., Viano, D. C. & Bir, C. A. Biomechanics of the head for Olympic boxer punches to the face. *Br. J. Sports Med.* **39**, 710–719 (2005).
31. Smith, M., Dyson, R., Hale, T. & Janaway, L. Development of a boxing dynamometer and its punch force discrimination efficacy. *J. Sports Sci.* **18**, 445–450 (2000).
32. Saltelli, A. et al. Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity index. *Comput. Phys. Commun.* **181**, 259–270 (2010).
33. Goodfellow, I., Bengio, Y. & Courville, A. *Deep Learning* (MIT Press, 2016).
34. Breiman, L. Random forests. *Mach. Learn.* **45**, 5–32 (2001).
35. Roozenbeek, B., Maas, A. I. & Menon, D. K. Changing patterns in the epidemiology of traumatic brain injury. *Nat. Rev. Neurol.* **9**, 231–236 (2013).
36. Morrison, B. et al. *A Tissue Level Tolerance Criterion for Living Brain Developed with an In Vitro Model of Traumatic Mechanical Loading* (SAE Technical Paper, 2003).
37. JE, W. & HR, L. Observations on the mechanism of brain concussion, contusion, and laceration. *Surg. Gynecol. Obstet.* **101**, 680–690 (1955).
38. Yoganandan, N. et al. Biomechanics of skull fracture. *J. Neurotrauma* **12**, 659–668 (1995).
39. Takhounts, E. G., Craig, M. J., Moorhouse, K., McFadden, J. & Hasija, V. *Development of Brain Injury Criteria (Bric)* (SAE Technical Paper, 2013).
40. Pearson, K. Note on regression and inheritance in the case of two parents. *Proc. R. Soc. Lond.* **58**, 240–242 (1895).
41. Hajiaghdam, M., Lan, I. S., Christian, C. W., Coats, B. & Margulies, S. S. Infant skull fracture risk for low height falls. *Int. J. Legal Med.* **133**, 847–862 (2019).
42. Wood, J. L. Dynamic response of human cranial bone. *J. Biomech.* **4**, 1–12 (1971).
43. Parvizi, J. & Damasio, A. Consciousness and the brainstem. *Cognition* **79**, 135–160 (2001).
44. Delano-Wood, L. et al. Brainstem white matter integrity is related to loss of consciousness and postconcussive symptomatology in veterans with chronic mild to moderate traumatic brain injury. *Brain Imaging Behav.* **9**, 500–512 (2015).
45. Parvizi, J. & Damasio, A. R. Neuroanatomical correlates of brainstem coma. *Brain* **126**, 1524–1536 (2003).
46. Sandhu, S. et al. Brainstem diffuse axonal injury and consciousness. *J. Clin. Imaging Sci.* **9**, 32 (2019).
47. Yoganandan, N., Pintar, F. A., Zhang, J. & Baisden, J. L. Physical properties of the human head: mass, center of gravity and moment of inertia. *J. Biomech.* **42**, 1177–1192 (2009).
48. Elliott, J. M. et al. Differential changes in muscle composition exist in traumatic and nontraumatic neck pain. *Spine* **39**, 39–47 (2014).
49. Edlow, B. L., Claassen, J., Schiff, N. D. & Greer, D. M. Recovery from disorders of consciousness: mechanisms, prognosis and emerging therapies. *Nat. Rev. Neurol.* **17**, 135–156 (2021).
50. Margulies, S. S. & Thibault, L. E. A proposed tolerance criterion for diffuse axonal injury in man. *J. Biomech.* **25**, 917–923 (1992).
51. Hernandez, F. et al. Six degree-of-freedom measurements of human mild traumatic brain injury. *Ann. Biomed. Eng.* **43**, 1918–1934 (2015).
52. Biecek, P. & Burzykowski, T. *Explanatory Model Analysis: Explore, Explain, and Examine Predictive Models* (CRC Press, 2021).
53. Haykin, S. *Neural Networks: A Comprehensive Foundation* (Prentice Hall PTR, 1998).

54. Bottou, L. *Proc. COMPSTAT'2010: 19th International Conference on Computational Statistics Paris France, August 22–27, 2010 Keynote, Invited and Contributed Papers*. 177–186 (Springer, 2010).
55. Henningsen, M. J. et al. Subject-specific finite element head models for skull fracture evaluation—a new tool in forensic pathology. *Int. J. Legal Med.* **138**, 1447–1458 (2024).

Acknowledgements

The authors would like to acknowledge funding from Home Office's Police STAR Fund. The authors would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work <https://doi.org/10.5281/zenodo.22558>.

Author contributions

Y.W. developed the numerical framework, analysed the data and wrote the paper. P.H. contributed to both coding and simulations. J.J. and M.J. provided biomechanical and medical guidance. J.O., S.B., and N.C. collected crime reports, analysed them and postprocessed them. J.M.P. provided guidance on machine learning algorithms. S.G. designed the overall research work. A.J. supervised the overall numerical work, designed the overall research, analysed the data and wrote the paper. All authors revised the paper.

Competing interests

The authors declare no competing interests.

Ethical approval

The studies involving the collection and analysis of criminal cases were reviewed and approved by the South Central-Berkshire Research Ethics Committee, the Thames Valley Police and National Crime Agency of UK. All documents were anonymised before postprocessing. This study adheres to the established standards of ethical research and inclusion. Access to the data was facilitated by Thames Valley Police and the National Crime Agency. Consent forms were obtained from the victims for all assault cases and were not required for deceased cases. The data used in this research were fully anonymised by the Thames Valley Police and the National Crime Agency before being shared with the academic authors for analysis.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s44172-025-00352-2>.

Correspondence and requests for materials should be addressed to Stan Gilmour or Antoine Jérusalem.

Peer review information *Communications Engineering* thanks Shaoju Wu and the other, anonymous, reviewer for their contribution to the peer review of this work. Primary Handling Editors: [Sandra Rugonyi, Saleem Denholme, Rosamund Daw].

Reprints and permissions information is available at <http://www.nature.com/reprints>

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

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025