Journal of Pediatric Surgery 54 (2019) 2353–2357

Contents lists available at ScienceDirect

Journal of Pediatric Surgery

journal homepage: www.elsevier.com/locate/jpedsurg

Trauma

Prediction of cervical spine injury in young pediatric patients: an optimal trees artificial intelligence approach☆

Dimitris Bertsimas [a], Peter T. Masiakos [b,*], Konstantinos S. Mylonas [b], Holly Wiberg [a]

[a] Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA, USA
[b] Division of Pediatric Surgery, Massachusetts General Hospital, Harvard Medical School, Boston, MA, USA

ARTICLE INFO

Article history:
Received 31 December 2018
Received in revised form 4 March 2019
Accepted 11 March 2019

Key words:
Cervical spine trauma
Cervical spine injury
Artificial intelligence
Machine learning
Optimal classification trees

ABSTRACT

Background: Cervical spine injuries (CSI) are a major concern in young pediatric trauma patients. The consequences of missed injuries and difficulties in injury clearance for non-verbal patients have led to a tendency to image young children. Imaging, particularly computed tomography (CT) scans, presents risks including radiation-induced carcinogenesis. In this study we leverage machine learning methods to develop highly accurate clinical decision rules to predict pediatric CSI.
Methods: The PEDSPINE I registry was used to investigate CSI in blunt trauma patients under the age of three. Predictive models were built using Optimal Classification Trees, a novel machine learning approach offering high accuracy and interpretability, as well as other widely used machine learning methods.
Results: The final Optimal Classification Trees model predicts injury based on overall Glasgow Coma Score (GCS) and patient age. This model has a sensitivity of 93.3% and specificity of 82.3% on the full dataset. It has comparable or superior performance to other machine learning methods as well as existing clinical decision rules.
Conclusions: This study developed a decision rule that achieves high injury identification while reducing unnecessary imaging. It demonstrates the value of machine learning in improving clinical decision protocols for pediatric trauma.
Type of study: Retrospective Prognosis Study.
Level of evidence: II

Data-driven protocols offer an evidence-based and objective approach to diagnosis and treatment decisions [1,2]. These advantages are especially relevant in the chaos and emotion of the emergency room. Trauma physicians treat a broad range of patient conditions and must make treatment decisions quickly, often based on limited information. The use of clinical decision rules, in tandem with a physician's individual instincts, improves consistency in this high-stress environment.

State-of-the-art machine learning methods, such as Artificial Neural Networks (ANN) and Gradient-Boosted Trees, have been introduced to a variety of clinical settings [3–6]. They have attracted much attention owing to their ability to parse through large quantities of data and capture complex trends between covariates in ways that outperform previously existing methods. However, these recently emergent methods suffer from a lack of interpretability, which is problematic in a clinical setting in which the rationale behind treatment decisions is critical. This has prompted debate as to the benefits and risks of employing such methods in the field [7,8].

The Optimal Classification Trees (OCT) algorithm addresses these issues through the construction of highly accurate and readily interpretable decision trees using an optimization approach, making it a strong fit for medical decision making [9]. This has been successfully applied in clinical settings, from oncology to pediatric surgery [10,11], and we now introduce it to the pediatric trauma space through the study of cervical spine injury (CSI).

Although the incidence of CSI in blunt trauma patients is low (<2%) [12], the risks of missed injuries lead physicians to be cautious in injury clearance and result in high reliance on imaging, including computed tomography (CT) [13]. Nevertheless, heavy dependence on imaging does come with risks of its own. Previous research has established the potential carcinogenic effects of radiation exposure from CT scans [14–17]. In addition, concerns about resource consumption, costs, and sedation risks have brought attention to the liberal use of imaging more broadly [18–20].

Several protocols have been developed for CSI clearance in both adult and pediatric settings [21–23], but these do not extend well to very young children who lack verbal cues and have different injury risk profiles. Our group constructed a multisite registry of over 12,000

☆ Funding: None.
* Corresponding author at: Harvard Medical School, Division of Pediatric Surgery, Massachusetts General Hospital, Warren 11, 55 Fruit Street, Boston, MA 02114, USA. Tel.: +1 617 680 5144.
E-mail address: pmasiakos@partners.org (P.T. Masiakos).

https://doi.org/10.1016/j.jpedsurg.2019.03.007
0022-3468/© 2019 Elsevier Inc. All rights reserved.

pediatric patients under the age of 3. Using this dataset we proposed the PEDSPINE criteria, a logistic-regression based model that determines CSI risk using Glasgow Coma Score, patient age, and mechanism of injury [12]. Anderson et al. [24] proposed another protocol for the same age group, but this relies on initial X-ray screening and provides no quantitative decision rule for injury clearance.

The aim of this post-hoc analysis of the PEDSPINE registry is to validate and refine existing protocols for young pediatric patients (0–3 years old) using modern machine learning methods. A cervical spine injury clearance decision rule is constructed using OCT, providing highly accurate and interpretable predictions that offer potential to reduce unnecessary imaging.

1. Methods

1.1. Study population

The PEDSPINE I database was queried for this study [12]. This international registry includes data on young pediatric patients (<3 years old) that were treated for blunt trauma in 22 trauma centers across the United States, Canada, and Brazil. Of these sites, 15 are pediatric level I trauma centers, six are level I adult trauma centers treating pediatric patients, and one is a level II adult trauma center. Deaths are omitted from the dataset since these records are generally incomplete and lack information on cervical spine injury.

1.2. Data and outcome measures

The PEDSPINE I registry includes several covariates for each trauma patient. Age and gender were reported for demographic features. In addition, the Glasgow Coma Score (separated into Eye, Verbal, and Motor subscores) and mechanism of injury (categorized as motor vehicle crash, fall, abuse, or other) were reported. Features that would not be known at the time of initial patient evaluation were excluded from the model, such as injury severity score, imaging results, and length of stay.

1.3. Machine learning analysis

1.3.1. Optimal trees

The OCT algorithm was used to develop the predictive model. OCTs build decision trees which optimize a specified loss function, with an additional complexity penalty to control the tree size. Each tree node is either a split node, which is characterized by a feature and threshold, or a leaf node, which is the terminal point for an observation and provides the predicted class label. The covariates and outcomes are given as inputs to the model, and the resultant model predicts a class label for each patient as injured or uninjured. This translates to imaging recommendations: patients predicted to be injured should be imaged, whereas those predicted to be uninjured statistically do not require imaging and in the clinical setting may be eligible to avoid an unnecessary procedure.

OCT improves upon CART, a tree-based method popularized in the 1980s, in that the decisions consider the tree's global structure while CART constructs a greedy decision tree [25]. The OCT algorithm is initialized with a greedy tree and then searches for improvements to the global objective by visiting nodes in a random order and considering various tree modifications. For leaf nodes, the algorithm considers the creation of a split at the leaf. For split nodes, it considers deleting the split or replacing it with a different split. The algorithm continues to consider nodes one-at-a-time until no more improvements are found. The full process is repeated with many random starts, and the tree with the lowest overall loss objective is selected as the final decision tree.

1.3.1.1. Model criterion. The OCT model was trained using the weighted misclassification rate as the loss function to be minimized. The standard misclassification rate is defined as the proportion of observations that are assigned the incorrect class label. This original metric is limited in that it does not reflect the fact that in our decision setting, false negatives are more problematic than false positives. In other words, it is worse to miss an injury than to misclassify someone as injured.

To reflect the priority toward injury identification, a weighted misclassification rate is used as the loss function. Suppose there are N patients, and the false positive (assigned “injured” label, but actually uninjured) and false negative (assigned “uninjured” label, but actually injured) counts are given as FP and FN, respectively. Given a weight w for missed injuries, the loss is computed as:

$$\text{loss} = \frac{\text{FP} + w * \text{FN}}{N}$$

The weight parameter represents the additional penalty on missed injuries. As the weight increases, the model favors more false positives than false negatives, resulting in more liberal injury predictions but fewer missed injuries.

1.3.2. Gradient-boosted trees

A gradient-boosted tree algorithm, introduced by Friedman, uses the ideas of boosting and gradient-descent to achieve higher predictive accuracy than traditional greedy tree-based methods [26]. Rather than fitting a single greedy tree and predicting a class label based on the observation's assigned leaf, this algorithm takes an iterative approach; at each stage, a new tree is trained to fit the prediction error that persists from previous trees, and the final predicted injury probabilities are made based on a linear combination of the individual trees. The probability threshold for labeling patients as injured was chosen to minimize the weighted misclassification criterion described in Section 1.3.1.1. While gradient-boosted trees have become popular in the machine learning community owing to their strong predictive performance, they lack interpretability and offer little visibility into the underlying mechanisms of predictions. The model in this paper was constructed using the *xgboost* package in R [27].

1.3.3. Logistic regression

A logistic regression model was also fit to provide a basis of comparison for model performance. As with gradient-boosted trees, the probability threshold for injury was chosen to minimize the weighted misclassification rate.

2. Results

2.1. Study population and characteristics

Our dataset included 9652 patients, of which 119 (1.23%) had a cervical spine injury and the remaining did not. Table 1 provides a summary of patient features by injury status. Seventy-eight (65.55%) of

Table 1
Comparison of patient features and imaging decisions by injury status.

		Uninjured (N = 9533)	Injured (N = 119)
Demographics	Age	1.3 (1.0–2.0)	2.0 (1.0–2.6)
	Male (%)	5578 (58.5%)	70 (58.8%)
GCS scores	GCS eye	4 (4–4)	1 (1–4)
	GCS verbal	5 (5–5)	1 (1–4)
	GCS motor	6 (6–6)	2 (1–4)
	GCS total	15 (15–15)	5 (3–11)
Injury type	Abuse (%)	555 (0.6%)	6 (0.5%)
	Fall (%)	3855 (40.4%)	24 (20.2%)
	MVA (%)	2396 (25.1%)	65 (54.6%)
	Other (%)	2727 (28.6%)	14 (11.8%)
Imaging Decisions	CT scan (%)	1722 (18.1%)	78 (65.5%)
	MRI (%)	203 (2.1%)	39 (32.8%)
	X-ray (%)	1577 (16.5%)	51 (42.9%)
	Overall (%)	2735 (28.7%)	102 (85.7%)

Reported values are medians with interquartile ranges or frequencies with percentages.

Table 2
Comparison of imaging rates by hospital type.

Hospital type	N (i)	Imaging rates			
		CT	MRI	X-ray	Overall
All Sites	9652 (119)	1800 (18.6%)	242 (2.5%)	1628 (16.9%)	2837 (29.4%)
Pediatric level I	5374 (63)	721 (13.4%)	134 (2.5%)	1330 (24.7%)	1561 (29.0%)
Adult level I	2691 (17)	561 (20.8%)	47 (1.7%)	152 (5.6%)	653 (24.3%)
Adult level II	1582 (39)	515 (32.6%)	61 (3.9%)	144 (9.1%)	619 (39.1%)

N: total number of patients; i: number of injured patients; CT: computed tomography; MRI: magnetic resonance imaging.

the injured patients received CT scans. Out of the 9533 uninjured patients in our dataset, 1722 (18.06%) received CT scans and 2736 (28.70%) received at least one form of imaging (CT, MRI, and x-rays). Table 2 shows a breakdown of the true imaging decisions for each hospital type, stratified by trauma level designation and pediatric setup (standalone children's hospital or general hospital).

2.2. Loss function calibration

The model fitting process began by calibrating the misclassification rate to appropriately bias the loss function toward the correct identification of injuries over noninjuries. The OCT algorithm was used to fit trees on a training dataset using various injury weights, ranging from 50 to 500. Fig. 1 illustrates the effect of injury weight on key performance metrics, evaluated on a separate test set. The key metrics of interest are sensitivity and specificity, which are indicative of the rates of missed injuries and missed noninjuries, respectively. The highest priority is achieving high sensitivity, but it is also important to maintain reasonable specificity to avoid unnecessary imaging. Negative predictive value (NPV) is also compared, but this is not a differentiating feature in assessing model quality because of the extremely low incidence of injury.

The sensitivity improves and specificity worsens as the injury weight increases: a higher penalty on false negatives results in fewer missed injuries but also more uninjured patients labeled as injured. The sensitivity does not increase beyond a weight of 300 because of overfitting on the training set as the rare injury events increasingly dominate the loss function calculation. In addition, the loss in specificity is relatively modest up to this weight and steeply declines afterward. Thus, models were trained using weighted misclassification with an injury weight of 300.

2.3. Machine learning models

2.3.1. Optimal trees

After determining the optimal injury weight, the training data were used to fit a final OCT model. The maximum tree depth and minimum bucket size were tuned through cross-validation, with the final chosen parameters as tree depth of 2 and minimum leaf size of 25. The resulting tree in Fig. 2 is simple and highly interpretable. The model predicts injury for patients with a total Glasgow Coma Score (GCS) of ≤13 or age above 2.58 years (approximately 31 months), and otherwise predicts no injury.

2.3.2. Gradient-boosted trees

The gradient-boosted trees algorithm returns a predicted injury probability for each patient. This is converted into a class prediction by choosing a threshold above which patients are labeled as injured. The threshold was chosen as .00439 since it is the minimizer of the weighted misclassification error. The model ranks GCS total score, GCS motor score, and patient age as the three most important variables, but no further interpretation of the prediction is available.

2.3.3. Logistic regression

The final logistic regression model predicts injury probability as a linear combination of the features. Higher injury probabilities are assigned to older children, females, fall or motor vehicle mechanisms of injury, and lower GCS scores. This method assumes an underlying linear structure and is unable to capture complexities in interactions between the variables.

All models were trained and validated using a random sample of 75% of the patients in this dataset. The remaining patients were set aside as a test group to evaluate the final models.

2.4. Outcomes

The outcomes on the full dataset are reported for all methods in Table 3. The OCT model reports a sensitivity of 93.28% on the full dataset: 111 of the 119 cervical spine injuries are correctly recovered in the model. It achieves a specificity of 82.34% and AUC of 90.43%. At the chosen injury thresholds, OCT obtains comparable sensitivity to the other two methods and significantly outperforms logistic regression in specificity. While it has a slightly lower specificity than gradient-boosted trees, the loss is marginal compared to OCT's gain in interpretability and delineation of an explicit decision rule.

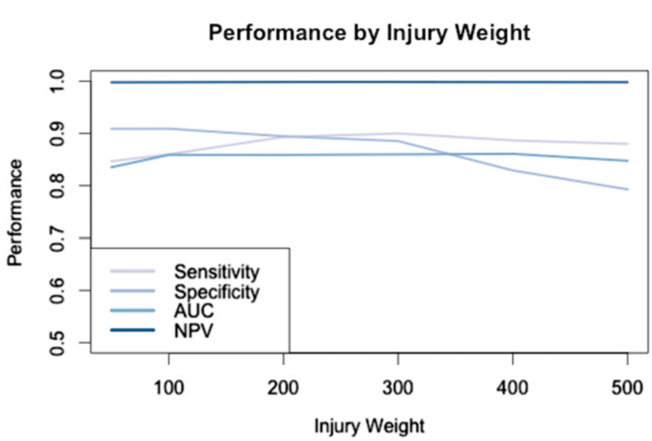

Fig. 1. Effect of loss function injury weight on predictive performance. AUC: area under the curve; NPV: negative predictive value.

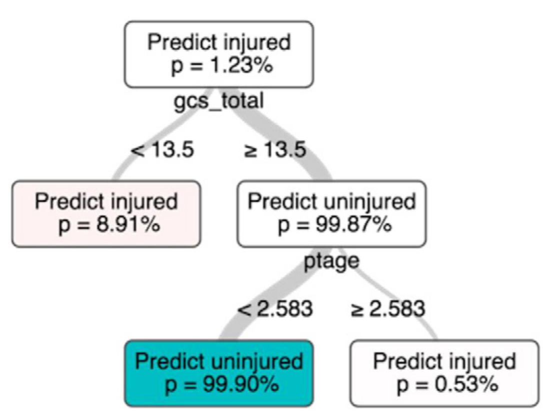

Fig. 2. Final Optimal Classification Tree.

Table 3
Comparison of machine learning methods on the full dataset.

	Sensitivity	Specificity	AUC
Optimal classification trees	93.28%	82.34%	90.43%
Gradient-boosted trees	94.96%	84.66%	96.69%
Logistic regression	95.8%	70.63%	94.06%

3. Discussion

The models developed in this paper provide an algorithmic approach to assessing cervical spine injury risk for patients under the age of three. Of the considered algorithms, OCT offers the best balance of performance and interpretability. The final OCT model predicts cervical spine injury incidence based on two basic features from a trauma examination, GCS score and patient age. The simplicity of the model suggests that patient outcomes can be predicted with high accuracy using only basic exam findings, even before considering the more granular information that might be available to a clinician. The high interpretability of the clinical decision rule makes it a viable tool for real-time use in trauma settings.

This model has strong performance in identifying cervical spine injuries while avoiding unnecessary imaging on uninjured patients. As reported above, the OCT model has a sensitivity of 93.28% and specificity of 82.34%, which translates to imaging for 17.66% of uninjured patients. The true CT scan decisions at the contributing sites yield a sensitivity of 65.55% and specificity of 81.94%, which translates to imaging for 18.06% of uninjured patients. The proposed model offers a slight advantage over the true physician decisions as measured by CT imaging rates for uninjured patients, and it has a clear edge in identifying injured patients who need imaging as measured by the sensitivity. When considering the combination of CT scans, MRIs, and x-rays, the true physician imaging decisions have a sensitivity of 85.71% and specificity of 71.31%. This indicates that 28.69% of uninjured patients received some sort of imaging. These results demonstrate that when considering imaging decisions more broadly, the OCT model offers a meaningful advantage in both sensitivity and specificity. Overall, the algorithmic approach achieves a notable reduction in overall imaging rates and comparable CT scan rates when compared to the true treatment decisions.

3.1. Comparison to other methods

The 2009 PEDSPINE study developed the only known cervical spine injury clearance protocol for patients younger than 3 years that does not rely on imaging [12]. In this model, a patient's risk score ranges from 0 to 8 and is calculated using four independent clinical predictors of CSI: GCS total score < 14 (3 points), GCS eye score = 1 (2 points), mechanism of injury = motor vehicle crash (2 points), and age = 2 years or older (1 point). Patients with a score of 0 or 1 are predicted to be uninjured, and higher scores are predicted to be injured and thus result in a recommendation of imaging. The PEDSPINE model is largely consistent with the tree presented in this paper: both identify an increased risk of injury for patients with GCS total scores lower than 13 as well as older patients. While these are the only factors in our model, PEDSPINE also assigns higher risk to patients with a low GCS eye subscore or motor vehicle accident (MVA) cause of injury.

An out-of-sample comparison of all models is presented in Fig. 3. The ROC curve is plotted for all candidate probability thresholds where applicable, and the symbols indicate the performance at the threshold chosen based on the weighted misclassification rate. Our model offers similar or superior performance in both of the key metrics. In particular, when considered against the PEDSPINE protocol, the OCT model yields comparable injury identification rates with a notable reduction in unnecessary imaging for uninjured patients. The proposed OCT model would save 1506 scans among the 9533 uninjured patients in the dataset. This is accomplished with a simpler model that relies on

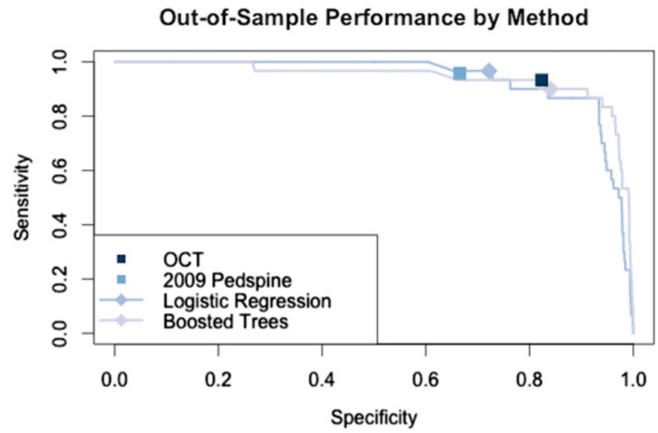

Fig. 3. Comparison of methods by sensitivity and specificity. OCT: Optimal Classification Trees.

fewer clinical features and attests to the power of machine learning to discern trends that may be obscured with traditional statistical approaches.

3.2. Clinical implications and recommendations

According to our analysis, patients with GCS > 13 and age less than 2.5 years have an exceptionally low risk of being injured. Using these criteria to clinically clear cervical spines, we primarily aim to reduce excessive CT scanning which is the most dangerous imaging modality in terms of ionizing radiation. However, reducing overall imaging is also important from a resource utilization and effective patient management standpoint. Indeed, it is well proven that plain films miss significant injuries and CT is a more appropriate and accurate test when there is true concern about the c-spine. Of note, level II adult trauma centers (39.1%) imaged nearly 10% more patients compared to level I pediatric centers (29.4%), Therefore, our guideline may have the largest impact in less specialized trauma centers.

Although the OCT model achieved impressive predictive performance, it is not a substitute for comprehensive clinical assessment. As previously acknowledged, 111 of the 119 cervical spine injuries were correctly recovered in the model. It is important to note that no statistical algorithm or any imaging test can reach 100% sensitivity; meaning that there is potential for missed injuries with any and all available support tools. In line with current clinical practice (using imaging), the machine learning approaches presented in this paper are meant to provide additional clinical support. In the case of equivocal clinical findings, decision making should default to clinical judgment.

3.3. Limitations

Our study has several limitations owing to its retrospective, multi-institutional, and multinational nature. First, we analyzed only basic clinical data, because the collection of more complex parameters was subject to wide variability among institutional trauma registries. Therefore, potentially useful information on hemodynamic presentation, associated injuries, detailed clinical signs and symptoms was not collected.

From a statistical standpoint, the infrequency of cervical spine injury occurrence in the dataset presents a challenge in developing machine learning approaches for this clinical setting. The model is at risk of overfitting the data and identifying injury indicators that are in fact simply characteristics of our injured patient sample and not generalizable to broader CSI clearance in the population. We used low depth trees and enforced a minimum sample size in each leaf to limit this issue.

Lastly, the data used in this study were collected from 1994 to 2004. Imaging protocols and child safety policies such as seatbelt laws have changed over the past 20 years, and thus there may be differences in the population of pediatric trauma patients with blunt head injury, as well as differences in injury detection and imaging protocols. In future work, we hope to conduct a similar analysis with an updated dataset using recent patient records to more accurately reflect the current setting of CSI clearance.

4. Conclusions

Overall, the OCT algorithm allows us to accurately identify understandable predictors of cervical spine injuries while imposing minimal structure on the model features and distribution. This nonparametric approach provides value in several ways. In cases such as this paper, it provides data-driven validation of relatively well-understood predictors of cervical spine injury. This results in more rigorous and evidence-based clinical decision protocols and lessens the reliance on varied physician experience and knowledge. Furthermore, this approach can identify novel predictors and interactions between features. This is particularly valuable in settings that lack medical consensus on diagnosis or treatment protocol.

The results from this paper demonstrate the potential for broader application of Optimal Trees in clinical decision recommendations. Through the construction of a readily interpretable and globally optimal decision tree, the OCT algorithm is well-suited for medical decision-making settings; it illustrates the power of developing methods that bridge the gap between interpretability and modern machine learning. In future work, we hope to extend this framework to other pediatric trauma settings.

References

[1] Amato F, López A, Peña-Méndez EM, et al. Artificial neural networks in medical diagnosis. J Appl Biomed 2013;11:47–58. https://doi.org/10.2478/v10136-012-0031-x.
[2] Miller DD, Brown EW. Artificial intelligence in medical practice: the question to the answer? Am J Med 2018;131:129–33. https://doi.org/10.1016/j.amjmed.2017.10.035.
[3] Khan J, Wei JS, Ringnér M, et al. Classification and diagnostic prediction of cancers using gene expression profiling and artificial neural networks. Nat Med 2001;7:673–9. https://doi.org/10.1038/89044.
[4] Greenspan H, Van Ginneken B, Summers RM. Guest editorial deep learning in medical imaging: overview and future promise of an exciting new technique. IEEE Trans Med Imaging 2016;35:1153–9. https://doi.org/10.1109/TMI.2016.2553401.
[5] Chen Y, Jia Z, Mercola D, et al. A gradient boosting algorithm for survival analysis via direct optimization of concordance index. Comput Math Methods Med 2013;2013. https://doi.org/10.1155/2013/873595.
[6] Esteva A, Kuprel B, Novoa RA, et al. Dermatologist-level classification of skin cancer with deep neural networks. Nature 2017;542:115–8. https://doi.org/10.1038/nature21056.
[7] Cabitza F, Rasoini R, Gensini GF. Unintended consequences of machine learning in medicine. JAMA 2017;318:517. https://doi.org/10.1001/jama.2017.7797.
[8] Tu JV. Advantages and disadvantages of using artificial neural networks versus logistic regression for predicting medical outcomes. J Clin Epidemiol 1996;49:1225–31. https://doi.org/10.1016/S0895-4356(96)00002-9.
[9] Bertsimas D, Dunn J. Optimal classification trees. Mach Learn 2017;106:1039–82. https://doi.org/10.1007/s10994-017-5633-9.
[10] Bertsimas D, Dunn J, Pawlowski C, et al. Applied informatics decision support tool for mortality predictions in patients with cancer. JCO Clin Cancer Inform 2018:1–11. https://doi.org/10.1200/CCI.18.00003.
[11] Bertsimas D, Dunn J, Velmahos GC, et al. Surgical risk is not linear. Ann Surg 2018; 268:574–83. https://doi.org/10.1097/SLA.0000000000002956.
[12] Pieretti-Vanmarcke R, Velmahos GC, Nance ML, et al. Clinical clearance of the cervical spine in blunt trauma patients younger than 3 years: a multi-center study of the American Association for the Surgery of Trauma. J Trauma 2009;67:543–50. https://doi.org/10.1097/TA.0b013e3181b57aa1.
[13] Davis JW, Phreaner DL, Hoyt DB, et al. The etiology of missed cervical spine injuries. J Trauma 1993;34:342–6. https://doi.org/10.1097/00005373-199303000-00006.
[14] Brenner DJ, Elliston CD, Hall EJ, et al. Estimated risks of radiation-induced fatal cancer from pediatric CT. Am J Roentgenol 2001;176:289–96. https://doi.org/10.2214/ajr.176.2.1760289.
[15] Pearce MS, Salotti JA, Little MP, et al. Radiation exposure from CT scans in childhood and subsequent risk of leukaemia and brain tumours: a retrospective cohort study. Lancet 2012;380:499–505. https://doi.org/10.1016/S0140-6736(12)60815-0.
[16] Brody AS, Frush DP, Huda W, et al. Radiation risk to children from computed tomography. Pediatrics 2007;120:677–82. https://doi.org/10.1542/peds.2007-1910.
[17] Miglioretti DL, Johnson E, Williams A, et al. The use of computed tomography in pediatrics and the associated radiation exposure and estimated cancer risk. JAMA Pediatr 2013;167:700–7. https://doi.org/10.1001/jamapediatrics.2013.311.
[18] Sanborn PA, Michna E, Zurakowski D, et al. Adverse cardiovascular and respiratory events during sedation of pediatric patients for imaging examinations. Radiology 2005;237:288–94. https://doi.org/10.1148/radiol.2371041415.
[19] Cote CJ, Notterman DA, Karl HW, et al. Adverse sedation events in pediatrics: a critical incident analysis of contributing factors. Pediatrics 2000;105:805–14. https://doi.org/10.1542/peds.105.4.805.
[20] Smith-Bindman R, Miglioretti DL, Larson EB. Rising use of diagnostic medical imaging in a large integrated health system. Health Aff 2008;27:1491–502. https://doi.org/10.1377/hlthaff.27.6.1491.
[21] Hoffman JR, Wolfson AB, Todd K, et al. Selective cervical spine radiography in blunt trauma: methodology of the national emergency X-radiography utilization study (NEXUS). Ann Emerg Med 1998;32:461–9. https://doi.org/10.1016/S0196-0644(98)70176-3.
[22] Stiell IG, Wells GA, Vandemheen KL, et al. The Canadian C-spine rule for radiography in alert and stable trauma patients. JAMA 2001;286:1841–8. https://doi.org/10.1001/jama.286.15.1841.
[23] Arbuthnot M, Mooney DP. The sensitivity and negative predictive value of a pediatric cervical spine clearance algorithm that minimizes computerized tomography. J Pediatr Surg 2017;52:130–5. https://doi.org/10.1016/j.jpedsurg.2016.10.031.
[24] Anderson RCE, Kan P, Vanaman M, et al. Utility of a cervical spine clearance protocol after trauma in children between 0 and 3 years of age. J Neurosurg Pediatr 2010;5:292–6. https://doi.org/10.3171/2009.10.PEDS09159.
[25] Breiman L, Friedman JH, Olshen RA, et al. Classification and regression trees; 1984; 358. https://doi.org/10.1201/9781315139470.
[26] Friedman JH. Greedy function approximation: a gradient boosting machine, vol. 29; 2001.
[27] Chen T, He T, Benesty M, et al. xgboost: Extreme Gradient Boosting; 2018.