Fully Automatic Deep Learning Approaches for Cervical Spine Fracture Characterisation in CT Images

Elena Goyanes*†, Carmen Lozano*†, Joaquim de Moura*†, Jorge Novo*† and Marcos Ortega*†
*VARPA Research Group, Instituto de Investigación Biomédica de A Coruña (INIBIC),
Universidade da Coruña, A Coruña, Spain
†Depart. of Computer Science and Information Technologies, Universidade da Coruña, A Coruña, Spain
Email: elena.goyanes.gonzalez@udc.es, c.lozano@udc.es joaquim.demoura@udc.es, jnovo@udc.es, mortega@udc.es

***Abstract*—Cervical spine fractures are a critical clinical condition that requires an accurate diagnosis to prevent severe neurological complications. Although Computed Tomography (CT) is the imaging modality of choice for cervical spine evaluation, manual interpretation takes time and is subject to inter-observer variability, motivating the development of robust automated solutions. In this work, we present a fully automatic deep learning pipeline for cervical spine fracture analysis on CT scans, the first in the state-of-the-art comprising three complementary tasks: fracture screening, vertebrae segmentation, and fracture localisation. Fracture screening is formulated as a classification problem to identify CT slices containing fractures. Vertebra segmentation is achieved through a transfer-learning strategy based on brain MRI pretraining, enabling effective transfer of knowledge between modalities to cervical CT. For fracture location, we explore two alternative approaches: coordinate-based regression and prediction of the bounding-box, providing precise spatial characterisation of fracture regions. The proposed pipeline is evaluated on large dataset comprised by 711,601 images from 2019 patients. The experimental results demonstrate strong and consistent performance across all three tasks, supporting accurate and efficient fracture characterisation. Overall, our methodology has the potential to assist radiologists by improving diagnostic accuracy while reducing analysis time and workload.**

***Index Terms*—Deep Learning, Fracture Screening, Vertebra Segmentation, Fracture localisation, Transfer Learning**

I. Introduction

Spine fractures are a major global health concern, especially in the cervical region, and are frequently related to traumatic events, ageing, and bone fragility [1]. Treatment for these fractures may be delayed or inappropriate due to the difficulty in the detection or misinterpretation as general back pain. Early and accurate identification of fractures is essential for preventing serious complications, improving patient outcomes, and helping clinicians to decide.

This work was supported by Government of Spain through the research projects with [grant number PID2023-148913OB-I00] and [grant number PID2024-161024OB-I00], funded by MICIU/AEI/10.13039/501100011033 and by ERDF, EU; Consellería de Educación, Universidade e Formación Profesional, Xunta de Galicia, Grupos de Referencia Competitiva, [grant number ED431C 2024/33], predoctoral grant [grant number ED481A-2023-152]. Also supported by the ISCIII under the grant [FORT23/00010] as part of the Programa FORTALECE of MICIU.

Computed Tomography (CT) has become a widely used imaging modality for spine evaluation, providing detailed cross-sectional images that allow clinicians to visualize complex structures with higher precision than conventional radiography [2]. Despite these benefits, manual CT scan analysis is time-consuming and depends on the experience of the radiologist. The need for automated support in fracture detection and localisation is highlighted by the possibility of diagnostic errors due to inter-observer variability.

Recent advances in deep learning, particularly convolutional neural networks (CNNs), have shown promising results for vertebral fracture analysis in medical imaging. The majority of studies focus on fracture classification, determining whether a vertebra is fractured or not [3]–[5]. These approaches often achieve high accuracy in identifying fractures, however, they provide limited information about the exact location or extent of the injury, which is very important for clinicians to make decisions and for treatment planning. A smaller subset of studies explores segmentation methods to delineate vertebral structures, providing more spatial context and facilitating the analysis of vertebral morphology. Few studies address precise fracture location, often using keypoint or landmark techniques [6]–[8]. However, these approaches are carried out using private datasets with small sample sizes, which restricts reproducibility, generalisation, and the establishment of standardised benchmarks. In 2025, Sutradhar *et al.* [4] proposed a comprehensive approach to automate fracture detection and classification of cervical spine injuries from CT scans. Consequently, to the best of our knowledge, there are no publicly available, large-scale studies that perform the three tasks: fracture screening, vertebrae segmentation using MRI transfer learning and spine fracture localisation in CT images, highlighting a significant gap and the need for automated methods capable of performing precise characterisation and quantitative assessment of vertebral fractures.

In this work, we present the first fully automatic deep learning framework for cervical spine fracture analysis in CT images, addressing three complementary tasks: fracture screening, vertebrae segmentation, and fracture localisation. Vertebra segmentation is performed using a transfer learn-

ing strategy from brain MRI data, enabling effective cross-modality knowledge transfer to cervical CT imaging. For fracture localisation, we explore both coordinate-based and bounding box prediction strategies to provide precise spatial information. The proposed framework is evaluated on a large, publicly available dataset, demonstrating its ability to support radiologists by improving diagnostic accuracy and efficiency while reducing clinical workload.

II. Materials and Methods

In this section, we describe the cervical spine CT dataset used in this work and the proposed fully automatic deep learning framework.

A. Dataset

The dataset used in this study is a subset of dataset from the RSNA Cervical Spine Fracture Detection Challenge [9], which provides high-resolution cervical spine CT scans acquired from multiple hospitals and clinical centres using state-of-the-art CT equipment. For the screening task, we used 711,601 images from 2019 patients, of which 52.4% were fractured patients. Regarding the segmentation task, we used images from 87 patients, which were the ones who had available vertebrae segmentation masks. For the localisation task, we focus on 7,218 2D slices from 961 patients with annotated fractures. Each fracture is defined by a bounding box using coordinates for the top-left corner and width and height for the extent of the region. This dataset captures a wide variability in fracture types and severity, providing a solid benchmark for training and evaluating automatic deep learning methods for the precise location of spine fractures.

B. Methodology

The overall workflow of the methodology is illustrated in Fig. 1. The proposed pipeline addresses three complementary tasks: fracture screening, vertebrae segmentation, and fracture localisation. The performance of the framework is evaluated for each task, with a particular focus on the precision and robustness of automated fracture screening, anatomical segmentation, and precise fracture localisation in CT images.

1) Spine fracture screening: The spine fracture screening task aims to automatically determine whether a CT slice contains a cervical spine fracture. This task serves as an initial filtering stage within the proposed framework, allowing subsequent processing steps to focus on clinically relevant slices. The problem is formulated as a binary classification task, where each input CT slice is classified as fractured or non-fractured.

To address this task, we evaluate convolutional neural network architectures: ResNet152 [10] and DenseNet161 [11]. Both models are initialised with weights pretrained on ImageNet to leverage large-scale natural image representations, ensuring faster convergence and improving feature generalisation in the medical imaging domain. Each network processes individual CT slices and outputs a probability score that indicates the presence of a fracture.

The screening models are trained using five independent repetitions to ensure robustness, with different splits of the dataset using a 60% for training, a 20% for validation and a 20% for test, and ensuring that all the images from a patient are on the same subset to avoid information leakage. Each repetition is trained for 30 epochs with a batch size of 4. Stochastic Gradient Descent (SGD) is employed as the optimiser, with an initial learning rate of 0.1, a momentum of 0.9, and a ReduceLROnPlateau scheduler that adaptively decreases the learning rate based on validation loss. The Binary Cross-Entropy (BCE) loss function is used to supervise the binary classification objective.

a) Evaluation: To assess the performance of the screening approach, we used the 5 classification metrics defined below, as well as Area Under the Receiver Operating Characteristic Curve (AUC), which measures discrimination capability independently of classification thresholds. Together, these metrics, that are commonly used in similar medical imaging classification tasks [12], offer a thorough evaluation of the effectiveness and reliability of the proposed screening models.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{1}$$

$$\text{Precision} = \frac{TP}{TP + FP} \tag{2}$$

$$\text{Recall} = \frac{TP}{TP + FN} \tag{3}$$

$$\text{F1} = 2 \cdot \frac{precision \cdot recall}{precision + recall} \tag{4}$$

$$\text{Specificity} = \frac{TN}{TN + FP} \tag{5}$$

2) Vertebrae segmentation: Vertebrae segmentation is performed to delineate cervical vertebral structures and provide anatomical context for subsequent fracture localisation. Accurate segmentation enables a better understanding of vertebral morphology and supports precise spatial interpretation of fracture locations.

For this task, we employ a convolutional neural network based on U-Net [13], which is adapted for cervical spine CT segmentation using a multimodal transfer learning strategy. A model pretrained on brain MRI segmentation is fine-tuned on CT data to leverage previously learnt anatomical features while adapting to differences in imaging modality, contrast, and structural appearance. This transfer learning approach facilitates robust vertebrae delineation despite the inherent heterogeneity between the MRI and CT domains.

To explore the impact of anatomy on segmentation performance, we evaluated two different labelling configurations: segmentation of cervical and thoracic vertebrae, comprising 14 vertebrae; and segmentation restricted to cervical vertebrae only, with 7 vertebrae masks. For each configuration, we conducted experiments both from scratch and using the model pretrained on MRI, enabling a systematic analysis of the advantages of transfer learning in this setting.

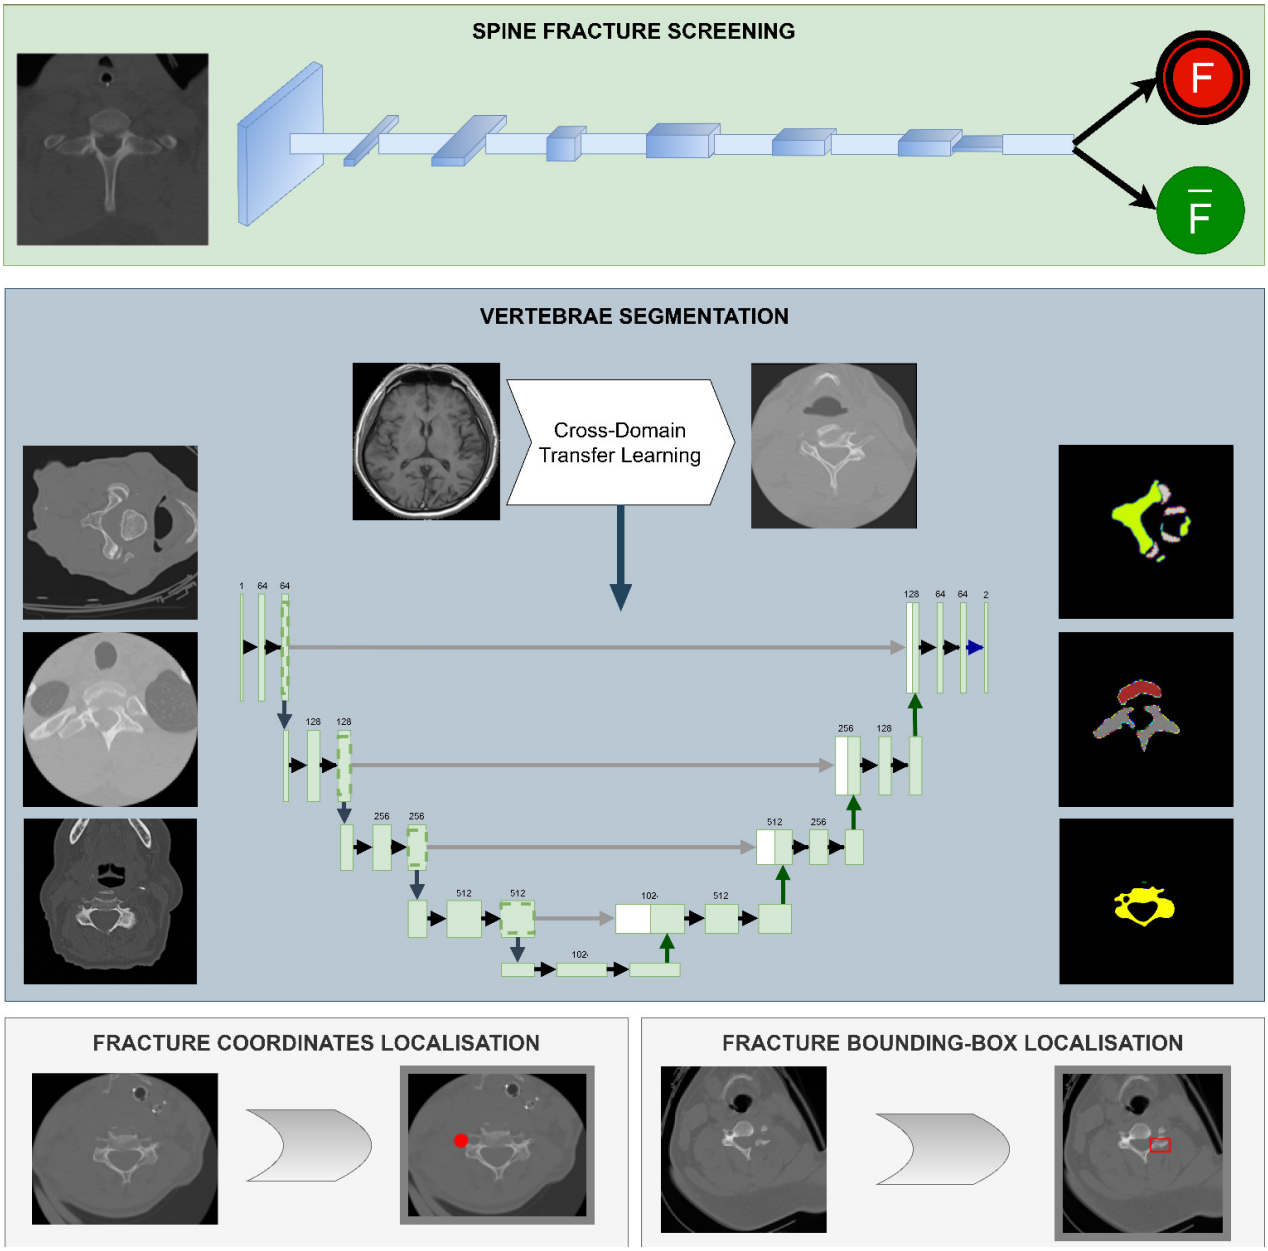

Fig. 1. Overview of the proposed methodology for the automatic characterisation of spine fractures. The workflow comprises three approaches: (1) spine fracture screening (top); (2) vertebra segmentation (middle); and (3) fracture localization (bottom), implemented through two complementary strategies: coordinate regression (left) and bounding-box prediction (right).

All segmentation models are trained for 300 epochs using a batch size of 16 and the Adam optimiser [14] with a learning rate of 0.001. Dice Loss is used to optimise the overlap between the predictions and the ground-truth masks. To ensure robustness and reproducibility, each experiment is repeated three times with random splits of the dataset, using 20% for testing and 80% of the data for training and validation, of which 85% was used for training and 15% for validation. The split was performed at the patient level to ensure that all images from a patient were on the same subset.

a) Evaluation: To quantitatively assess vertebra segmentation performance, we employed several complementary metrics widely used in similar medical image segmentation tasks [15]. Specifically, we analyse Accuracy (Eq. 1), Precision (Eq. 2), Recall (Eq. 3), and F1-score (Eq. 4), together with the Jaccard Index (Eq. 6), which measures the overlap between the predicted and ground-truth masks.

$$\text{Jaccard} = \frac{TP}{TP + FP + FN} \tag{6}$$

3) Spine fracture localisation: To evaluate the proposed automatic framework for spine fracture location in CT images, we considered two alternative tasks: predicting a single point representing the fracture position and estimating a bounding

box around the fracture region. Coordinate regression offers a precise location, while bounding boxes offer interpretable regions for clinicians. For both tasks, we assessed four representative deep convolutional architectures: DenseNet-121, DenseNet-161 [11], ResNet-50, and ResNet-101 [10], selected due to their proven effectiveness in medical image analysis and location.

All networks were independently trained following the same protocol. Models were initialised with weights pretrained on ImageNet and optimised using SGD with a learning rate of 0.001, a momentum of 0.9, and batch size of 16. The Mean Squared Error (MSE) was used as the loss function for both coordinate and bounding-box regression. All ground-truth coordinates and bounding-box values are expressed in pixel units. The data set was partitioned into 20% for testing and 80% of the data for training and validation, of which 85% was used for training and 15% for validation, ensuring that all images from the same patient were assigned to the same split. A ReduceLROnPlateau scheduler was applied to decrease the learning rate when validation performance becomes stagnant, and early stopping was employed with patience of 15 epochs to prevent overfitting. To ensure statistical robustness, each experiment was repeated 5 times, using a maximum of 300 training epochs per run.

a) Evaluation: To assess the precision of the proposed localisation approaches, we used three standard regression metrics [16], described below.

- **Mean Squared Error (MSE):** measures the average squared difference between the predicted and ground-truth values. Lower MSE indicates better predictive performance, although the metric is sensitive to outliers.

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2 \tag{7}$$

- **Root Mean Squared Error (RMSE):** corresponds to the square root of the MSE and provides an interpretable estimate of the average prediction error in the same units as the target variable.

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{8}$$

- **Mean Absolute Error (MAE):** computes the average absolute difference between the predicted and ground-truth values. It is easier to interpret than MSE and less sensitive to outliers.

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \tag{9}$$

III. Experimental results

This section presents the experimental results obtained for the three components of the proposed pipeline: fracture screening, vertebrae segmentation, and spine fracture localisation. For each task, we report quantitative performance metrics, compare different model architectures, and analyse the impact of training strategies such as transfer learning and network depth. Representative qualitative examples are also provided to complement the numerical evaluation.

A. Spine fracture screening

Table I summarizes the performance of the two evaluated architectures for the fracture screening task. Both DenseNet-161 and ResNet-152 achieve consistently high values across all metrics, reflecting their strong ability to discriminate between fractured and non-fractured slices. Performance is reported in terms of Accuracy, Precision, Recall, F1-score, AUC, and Specificity.

ResNet-152 achieves the best overall performance, obtaining the highest Accuracy (0.9943 $\pm$ 0.0014), F1-score (0.9946 $\pm$ 0.0013), and AUC (0.9943 $\pm$ 0.0014). DenseNet-161 also performs competitively, with only marginally lower values across all metrics. These results demonstrate that deep convolutional architectures pretrained on ImageNet are effective for CT slice classification, enabling reliable detection of patterns related to fractures.

B. Vertebrae segmentation

Quantitative segmentation results for the two anatomical labelling approaches: cervical and combined cervical and thoracic vertebrae segmentation are presented in Table II. Performance is evaluated using the metrics: Accuracy, Precision, Recall, F1-score, and the Jaccard Index.

For the cervical approach, the transfer learning strategy achieves the highest performance, reaching an Accuracy of 0.9976 $\pm$ 2.98e–5 and a Jaccard Index of 0.9954 $\pm$ 0.0081. Similar behaviour is observed when segmenting both cervical and thoracic vertebrae, where the pretrained model again outperforms training from scratch, achieving an Accuracy of 0.9964 $\pm$ 3.10e–5 and a Jaccard score of 0.9930 $\pm$ 0.0081. These results confirm that cross-modality transfer learning from brain MRI provides a substantial improvement, enhancing convergence stability and enhancing segmentation quality across vertebral classes.

Overall, both segmentation approaches present high and consistent results, validating their suitability for reliable vertebrae delineation in cervical CT imaging.

C. Spine fracture localisation

To evaluate the performance of the proposed localisation models, we use the regression metrics previously explained. Tables III and IV report the quantitative results for the two localisation strategies: coordinate regression and bounding-box prediction. Both strategies are evaluated using DenseNet and ResNet architectures of varying depths.

For fracture coordinate regression (Table III), DenseNet-161 achieves the best performance across all metrics, obtaining the lowest MSE (1214.64 $\pm$ 51.65), RMSE (34.81 $\pm$ 0.74), and MAE (26.81 $\pm$ 0.42). A similar trend is observed for the bounding-box localisation task (Table IV), in which DenseNet-161 again outperforms the remaining architectures, with the lowest MSE (1202.33 $\pm$ 64.46), RMSE (34.54 $\pm$ 0.94), and MAE (26.51 $\pm$ 0.66).

TABLE I

FRACTURE SCREENING RESULTS IN TERMS OF MEAN AND STANDARD DEVIATION. VALUES ARE REPORTED AS MEAN ± STANDARD DEVIATION.

Architecture	**Accuracy** (↑)	**Precision** (↑)	**Recall** (↑)	**F1** (↑)	**AUC** (↑)	**Specificity** (↑)
DenseNet-161	0.9930 ± 0.0010	0.9951 ± 0.0008	0.9918 ± 0.0013	0.9934 ± 0.0009	0.9931 ± 0.0010	0.9944 ± 0.0010
ResNet-152	**0.9943 ± 0.0014**	**0.9954 ± 0.0011**	**0.9939 ± 0.0016**	**0.9946 ± 0.0013**	**0.9943 ± 0.0014**	**0.9948 ± 0.0012**

TABLE II

SEGMENTATION RESULTS FOR EACH APPROACH: CERVICAL VERTEBRAE SEGMENTATION AND CERVICAL AND THORACIC VERTEBRAE SEGMENTATION. VALUES ARE REPORTED AS MEAN ± STANDARD DEVIATION.

Approach	**Transfer Learning**	**Accuracy** (↑)	**Jaccard** (↑)	**F1** (↑)	**Precision** (↑)	**Recall** (↑)
Cervical	✗	0.9944 ± 0.0069	0.9890 ± 0.0136	0.9944 ± 0.0069	0.9944 ± 0.0069	0.9944 ± 0.0069
	✓	**0.9976 ± 2.98e-5**	**0.9954 ± 0.0081**	**0.9977 ± 0.0000**	**0.9977 ± 0.0000**	**0.9977 ± 0.0000**
Cervical and thoracic	✗	0.9890 ± 0.0105	0.9785 ± 0.2045	0.9890 ± 0.1055	0.9880 ± 0.1045	0.9870 ± 0.1044
	✓	**0.9964 ± 3.10e-5**	**0.9930 ± 0.0081**	**0.9965 ± 1.42e-14**	**0.9965 ± 1.42e-14**	**0.9965 ± 1.42e-14**

TABLE III

RESULTS FOR SPINE FRACTURE COORDINATES LOCALISATION. VALUES ARE REPORTED AS MEAN ± STANDARD DEVIATION.

Architecture	**MSE** (↓)	**RMSE** (↓)	**MAE** (↓)
DenseNet-121	1480.91 ± 47.66	38.47 ± 0.62	29.75 ± 0.51
DenseNet-161	**1214.64 ± 51.65**	**34.81 ± 0.74**	**26.81 ± 0.42**
ResNet-50	1398.52 ± 121.41	37.34 ± 1.64	28.68 ± 1.15
ResNet-101	1509.32 ± 474.03	38.44 ± 5.56	29.09 ± 3.25

TABLE IV

RESULTS FOR SPINE FRACTURE BOUNDING BOX LOCALISATION. VALUES ARE REPORTED AS MEAN ± STANDARD DEVIATION.

Architecture	**MSE** (↓)	**RMSE** (↓)	**MAE** (↓)
DenseNet-121	1484.82 ± 43.52	38.41 ± 0.56	29.75 ± 0.45
DenseNet-161	**1202.33 ± 64.46**	**34.54 ± 0.94**	**26.51 ± 0.66**
ResNet-50	1372.30 ± 255.63	36.75 ± 3.24	28.32 ± 2.62
ResNet-101	1276.79 ± 76.76	35.60 ± 1.07	27.19 ± 0.64

These results indicate that deeper DenseNet architectures provide superior feature extraction capabilities for localisation, outperforming all tested ResNet variants. Moreover, the consistency across both localisation strategies demonstrates the robustness of DenseNet-161 in capturing structural features associated with cervical spine fractures.

D. Qualitative analysis

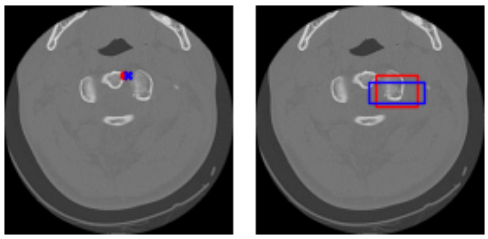

Fig. 2. Visual comparison between ground-truth (red) and predicted (blue) for each localisation approach: spine fracture location coordinates (left) and spine fracture bounding-box (right).

In addition to the quantitative analysis, Figure 3 provides representative qualitative examples for the segmentation task. The qualitative results are consistent with the strong quantitative performance reported in Table II, further confirming that the predictions produced by the models are coherent and of high visual quality.

In Figure 2, we show a visual comparison between the ground-truth annotations and the predictions obtained with both localisation strategies: coordinate regression (left) and bounding-box prediction (right). These examples illustrate how the proposed models accurately identify the fracture region. In particular, the close alignment between predicted and ground-truth annotations demonstrates the effectiveness of the proposed pipeline in capturing anatomical details relevant to cervical spine fracture location, further supporting the quantitative results reported in Tables III and IV.

IV. CONCLUSION

In this work, we presented the first fully automatic and comprehensive pipeline for cervical spine fracture analysis in CT images, covering three key tasks: fracture screening, vertebrae segmentation, and fracture localisation. The proposed framework was evaluated using a large, publicly available CT dataset. Across all tasks, our experimental results demonstrate that the deep learning architectures tested can effectively extract clinically relevant information for cervical spine fracture characterisation.

For the fracture screening task, both DenseNet-161 and ResNet-152 achieved high performance across all classification metrics, with ResNet-152 obtaining the best overall results. These findings highlight the capability of deep convolutional models pretrained on large-scale natural image datasets to accurately discriminate between fractured and non-fractured CT images.

In the vertebrae segmentation task, we investigated two anatomical approaches: only cervical and combined cervical and thoracic segmentation, trained both from scratch and with cross-modality transfer learning from brain MRI. In both settings, the transfer learning strategy clearly outperformed training from scratch, with the best results obtained for the cervical configuration, as expected. These experiments confirm the effectiveness of leveraging pretrained anatomical features to improve segmentation robustness in CT imaging.

Finally, for fracture location, we explored two complementary approaches: coordinate regression and bounding-box pre-

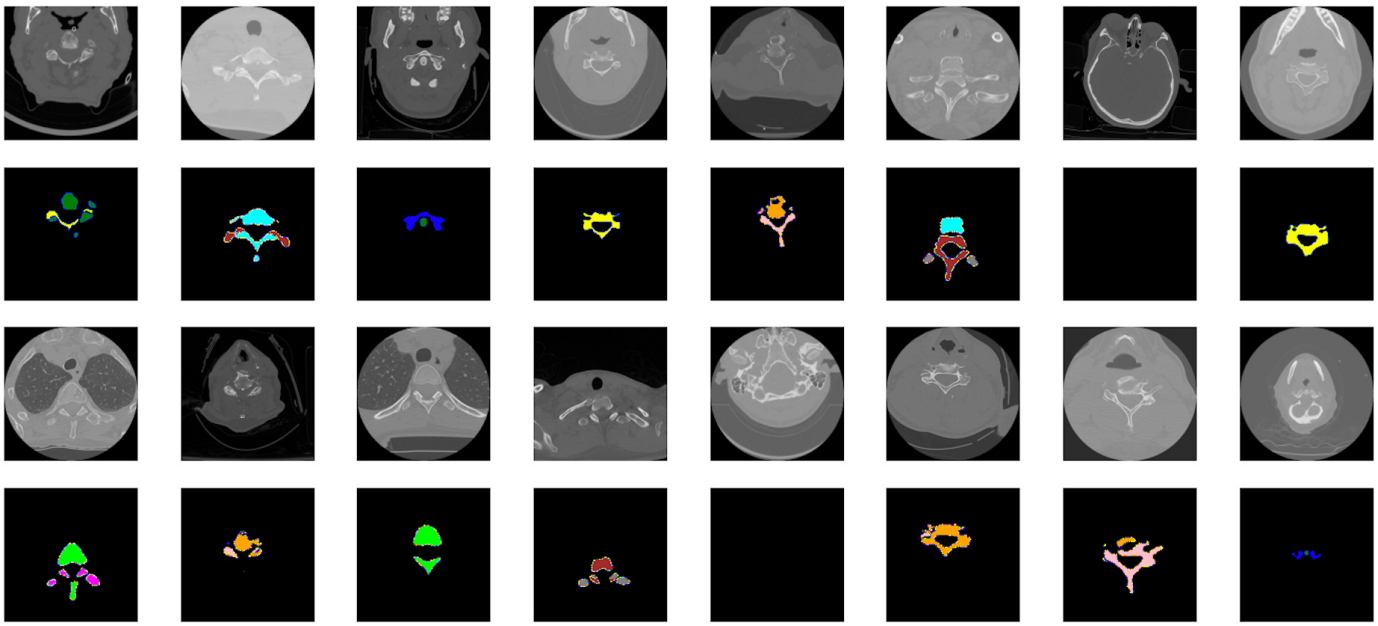

Fig. 3. Qualitative examples of input CT slices (top) and corresponding segmentation outputs (bottom). Color legend used for vertebra segmentation: Background, Vertebra 1, Vertebra 2, Vertebra 3, Vertebra 4, Vertebra 5, Vertebra 6, Vertebra 7, Vertebra 8, Vertebra 9, Vertebra 10, Vertebra 11.

diction. Across both methods and all evaluated architectures, DenseNet-161 consistently achieved the lowest localisation errors, demonstrating superior feature extraction capability for prediction. Qualitative visualisations further validated the quantitative findings, supporting the reliability of the proposed localisation strategies.

Overall, this study provides the first unified evaluation of cervical spine fracture screening, segmentation, and localisation on a large public CT dataset, establishing a benchmark for future research in automated spine fracture analysis.

Future work will include extending the methodology to additional imaging modalities such as radiography, as well as incorporating a dedicated 3D visualisation module to enhance interpretability and clinical applicability.

References

[1] T. Tsuda, "Epidemiology of fragility fractures and fall prevention in the elderly: a systematic review of the literature," *Current orthopaedic practice*, vol. 28, no. 6, pp. 580–585, 2017.

[2] W. A. Kalender, *Computed tomography: fundamentals, system technology, image quality, applications*. John Wiley & Sons, 2011.

[3] P. Chłąd and M. R. Ogiela, "Deep learning and cloud-based computation for cervical spine fracture detection system," *Electronics*, vol. 12, no. 9, p. 2056, 2023.

[4] D. Sutradhar, N. M. Fahad, M. A. K. Raiaan, M. Jonkman, and S. Azam, "Cervical spine fracture detection utilizing yolov8 and deep attention-based vertebrae classification ensuring xai," *Biomedical Signal Processing and Control*, vol. 101, p. 107228, 2025.

[5] M. Singh, U. Tripathi, K. K. Patel, K. Mohit, and S. Pathak, "An efficient deep learning based approach for automated identification of cervical vertebrae fracture as a clinical support aid," *Scientific Reports*, vol. 15, no. 1, p. 25651, 2025.

[6] B. Glocker, D. Zikic, E. Konukoglu, D. R. Haynor, and A. Criminisi, "Vertebrae localization in pathological spine ct via dense classification from sparse annotations," in *International conference on medical image computing and computer-assisted intervention*. Springer, 2013, pp. 262–270.

[7] M. Pisov, V. Kondratenko, A. Zakharov, A. Petraikin, V. Gombolevskiy, S. Morozov, and M. Belyaev, "Keypoints localization for joint vertebra detection and fracture severity quantification," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2020, pp. 723–732.

[8] A. Zakharov, M. Pisov, A. Bukharaev, A. Petraikin, S. Morozov, V. Gombolevskiy, and M. Belyaev, "Interpretable vertebral fracture quantification via anchor-free landmarks localization," *Medical Image Analysis*, vol. 83, p. 102646, 2023.

[9] H. M. Lin, E. Colak, T. Richards, F. C. Kitamura, L. M. Prevedello, J. Talbott, R. L. Ball, E. Gumeler, K. W. Yeom, M. Hamghalam *et al.*, "The rsna cervical spine fracture ct dataset," *Radiology: Artificial Intelligence*, vol. 5, no. 5, p. e230034, 2023.

[10] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

[11] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 4700–4708.

[12] E. Goyanes, J. De Moura, A. Lelong, P. Robles-Amor, J. M. Martínez-De-La-Casa, J. Moreno-Montañes, F. J. Muñoz-Negrete, I. Uña, J. Novo, and M. Ortega, "Domain-specific multi-transfer approaches for deep learning-based glaucoma screening in high myopia patients," *Expert Systems with Applications*, p. 129265, 2025.

[13] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.

[14] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[15] E. Goyanes, J. de Moura, J. Novo, J. I. Fernández-Vigo, J. Á. Fernández-Vigo, and M. Ortega, "Fully-automatic segmentation of the ciliary muscle using anterior segment optical coherence tomography images," in *2022 International Joint Conference on Neural Networks (IJCNN)*. Ieee, 2022, pp. 1–8.

[16] A. V. Tatachar, "Comparative assessment of regression models based on model evaluation metrics," *International Research Journal of Engineering and Technology (IRJET)*, vol. 8, no. 09, pp. 2395–0056, 2021.