

Explainable vertebral fracture analysis with uncertainty estimation using differentiable rule-based classification

Victor Wählstrand Skärström¹[0000-0001-6569-120X], Lisa Johansson²[0009-0005-6005-3838], Jennifer Alvé¹[0000-0003-4195-9325], Mattias Lorentzon²[0000-0003-0749-1431], and Ida Häggström^{1,2}[0000-0001-9178-6683]

¹ Chalmers University of Technology, SE-41296 Göteborg, Sweden

² Sahlgrenska Academy at University of Gothenburg, SE-40530 Göteborg, Sweden

Abstract

We present a novel method for explainable vertebral fracture assessment (XVFA) in low-dose radiographs using deep neural networks, incorporating vertebra detection and keypoint localization with uncertainty estimates. We incorporate Genant's semi-quantitative criteria as a differentiable rule-based means of classifying both vertebra fracture grade and morphology. Unlike previous work, XVFA provides explainable classifications relatable to current clinical methodology, as well as uncertainty estimations, while at the same time surpassing state-of-the-art methods with a vertebra-level sensitivity of 93% and end-to-end AUC of 97% in a challenging setting. Moreover, we compare intra-reader agreement with model uncertainty estimates, with model reliability on par with human annotators.

Keywords: Vertebral fracture assessment · detection · uncertainty quantification · explainability · rule-based classification · compression · morphology

1 Introduction

Vertebral compression fractures, osteoporotic deformations due to the loss of bone mass in the spine, are a common and serious complication of osteoporosis and a leading cause of impaired mobility and increased mortality [15]. Vertebral fracture assessment (VFA) using X-ray scans is a critical screening method used to detect osteoporotic fractures and apply interventional measures early [10]. The number and location of fractures is an important indicator of disease progression, but VFA is challenging, requiring time-consuming annotation of vertebrae by expert clinicians for reliable diagnoses. Patients are often elderly with difficulty holding position, and scans are commonly low-dose, decreasing image quality.

Genant's semi-quantitative method (GSQ) is a common means of VFA, where vertebrae in a spinal X-ray scan are annotated with six keypoints, and fractures are classified per morphology and severity [7], see Figure 1. In the semi-quantitative approach, readers make a compound judgement based on keypoint positions, visual features and adjacent vertebrae. The complexity of VFA makes




Fig. 1: *Visualization of GSQ*. Vertebral fracture morphology in Genant’s system, with landmarks $\mathbf{x}$ (dots). The posterior, middle, and anterior heights (lines) are used to determine deformation morphology and severity. Fig. adapted from [7].

the adoption of automated methods slow, and it is inhibited by poor understanding by clinicians, and low reliability and sensitivity of existing methods.

Previous work on automated VFA has focused on image level classification of the presence of vertebrae, due to the expensive nature of annotation. However, research indicates that the number and severity of compressions is highly predictive of future fractures.[14] Most authors employ convolutional neural networks (CNNs) for grade classification [18,6,17], using Grad-CAM [21] for weak localization and model interpretability. None of these methods provide direct localization of vertebrae, and methods like CAM and saliency maps do not explain model reasoning, only visually ranking feature contributions to predictions without providing tractable reasoning on their use in the model. Deep learning methods are thus commonly criticized for their “black box” classification and poor explainability, reducing trust and inhibiting the adaptation of methods in healthcare. Various approaches have addressed interpretability post-hoc for neural networks, as surveyed by e.g. Kouchaki et al. [12], while *interpretability* may be defined as an understanding of how an underlying model works, *explainability* reflects how the model came up with a given result. [9] Unlike neural networks, decision trees are intrinsically explainable, but typically not differentiable.

The localization of vertebrae has been studied in various imaging modalities, e.g. planar X-ray and computed tomography, e.g. using CNNs and graph neural networks [1,22]. For VFA, localization is not enough, and lack of annotated vertebrae make detection approaches rare. Recent end-to-end contributions include Cheng et al. [3], presenting a method of classification of trauma-induced fractures, where segmented vertebra features such as ridge heights were used as input to a random forest, yielding high sensitivity. Like previous work, their method operates on fairly high-contrast radiographs and does not differentiate fracture severity or relate their classifications to existing clinical methodology. Beside explainability, previous work rarely address model trust, an important factor in high stakes applications such as healthcare. An aspect of this is model uncertainty, for example confidence or error intervals imposing limits on model applicability. Dong et al. [5] use CNN ensembling for prediction, showing reliability across datasets. Residual log-likelihood [13] introduces a fast and light-weight alternative to expensive ensembling for regression tasks, using a normalizing flow [20] to model the error distribution of regressed points.




Fig. 2: *Model overview.* A detector $d_{\theta'}$ is trained to estimate bounding boxes (red) from vertebra in a spinal image (green). Ground truth crops with keypoints (green) are separately used to train a keypoint regressor $f_{x, \theta}$, subsequent vertebra classifiers $f_{img, \theta}, f_{kps}$, as well as estimate the keypoint distribution (red points).

1.1 Contributions

To our knowledge, no previous work has addressed VFA in low-dose, high-noise radiographs by incorporating GSQ directly into neural network classification. Our main contributions are creating a novel, explainable and reliable detection model of vertebral fractures through *i)* state-of-the-art detectors for vertebra localization, *ii)* estimation of vertebra keypoints with uncertainty quantification and *iii)* an innovative classification scheme using estimated vertebral keypoints to compute rule-based objective from GSQ. For this purpose we implement a differentiable decision tree, which enables us to propagate the keypoint uncertainty to the final fracture classification.

2 Methodology

An overview of XVFA is shown in Figure 2. We employ a two-stage method of coarse bounding box detection of vertebrae using a state-of-the-art (SOTA) detector, and regress vertebral keypoints and estimate of the empirical model error distribution using normalizing flows. Vertebrae are subsequently classified using a combination of GSQ criteria and a neural image classifier.

2.1 Coarse bounding box detection

In this paper, we use a SOTA detection transformer architecture, DETR with Improved DeNoising Anchor Boxes (DINO) with a ResNet50 backbone for localization [24, 8], having lower footprint and higher accuracy than its progenitor, DETR [2]. We use the same loss $\mathcal{L}_{det}$ as defined in DINO and DETR,

$$\mathcal{L}_{det} = \sum_i -\log \hat{p}(c_i) + \mathbf{I}_{c_i \neq \emptyset} \left(\lambda_{iou} \mathcal{L}_{iou}(b_i, \hat{b}_i) + \lambda_{l1} \|b_i - \hat{b}_i\|_1 \right), \quad (1)$$

computed between ground truth bounding boxes b and predicted boxes $\hat{b}$ matched with the Hungarian algorithm and regressed with generalized IoU loss $\mathcal{L}_{iou}$, [19] and ℓ_1 -loss. $\hat{p}(c)$ denotes the predicted class of the matched predicted object. If the matched object is predicted to belong to a non-object class, it is omitted from the loss, denoted with the indicator function, $\mathbf{I}_{c_i \neq \emptyset}$. In our case, we predict the presence or absence of vertebrae, corresponding to two classes in DETR.

2.2 Fine vertebra detection

Genant’s Semi-Quantitative Method Consider vertebra keypoints $\mathbf{x}_{ij} = (x_{ij}, y_{ij})$, as defined in Figure 1. The morphology of fractures varies between *normal*, *concave*, *wedge* and *crush* deformities, quantified in terms of vertebral heights h between the posterior, middle or anterior parts of the ridges:

$$h_p = \|\mathbf{x}_{up} - \mathbf{x}_{lp}\|, \quad h_m = \|\mathbf{x}_{um} - \mathbf{x}_{lm}\|, \quad h_a = \|\mathbf{x}_{ua} - \mathbf{x}_{la}\|, \quad (2)$$

where $\|\cdot\|$ denotes the 2-norm. Fracture morphology may be characterized through the anterior-posterior, middle-posterior and middle-anterior ratios between heights,

$$\text{APR} = \frac{h_a}{h_p}, \quad \text{MPR} = \frac{h_m}{h_p}, \quad \text{MAR} = \frac{h_m}{h_a}. \quad (3)$$

A normal vertebra is expected to be approximately rectangular and have equal height ratios of 1, within a tolerance 20%. A wedge deformity has $\text{MPR} < 1$ and $\text{MAR} > 1$, and inversely for crush fractures. Concave fractures have $\text{MPR} < 1$ and $\text{MAR} < 1$. The inverse case of convex vertebrae are considered normal.

According to GSQ, a mild fracture is defined as a relative ratio loss of 20-25% of vertebral height, a moderate fracture as 25-40% loss, and a severe fracture as more than 40% loss. Due to the small differences between classes, they are typically grouped as *normal+mild* and *moderate+severe*.

Keypoint estimation using residual log-likelihood Using a ResNet18 [8], vertebra images are encoded as features $\mathbf{z}$. As per residual log-likelihood estimation (RLE) [13], we estimate the log-likelihood $\log p_{\theta, \phi}(\mathbf{x}|\mathbf{z})$ of vertebra keypoints $\mathbf{x}$ conditioned on the visual features $\mathbf{z}$ by modelling the keypoint error distribution with a normalizing flow, f_ϕ , in our case RealNVP [4]. Like in maximum likelihood estimation, we assume an approximate target distribution, e.g. Gaussian or Laplace, where a trainable keypoint regressor predicts the centre and scale $f_\theta(\mathbf{z}) = (\hat{\boldsymbol{\mu}}, \hat{\mathbf{b}})$ of the distribution, minimizing the negative log-likelihood $-\log p_{\theta, \phi}(\mathbf{x}|\mathbf{z})$. The normalizing flow maps an initial distribution to a flow distribution $\mathbf{x}' \sim p_\phi$ such that the final distribution of keypoints $\mathbf{x}|\mathbf{z} \sim p_{\theta, \phi}$ is given by shifting and scaling, $\mathbf{x} = \hat{\mathbf{b}} \mathbf{x}' + \hat{\boldsymbol{\mu}}$. The log-likelihood can thus be written

$$\log p_{\theta, \phi}(\mathbf{x}|\mathbf{z}) = \log p_\phi \left(\frac{\mathbf{x} - \hat{\boldsymbol{\mu}}}{\hat{\mathbf{b}}} \right) - \log \hat{\mathbf{b}}. \quad (4)$$

Specifically, we assume that the predicted normalizing flow decomposes into a base distribution $q(\mathbf{x}')$ (e.g. a Gaussian) and a residual distribution $g_\phi(\mathbf{x}')$ up to a constant factor, such that $p_\phi = q + g_\phi$, yielding

$$\mathcal{L}_{rle} = \sum_i -\log q\left(\frac{\mathbf{x}_i - \hat{\boldsymbol{\mu}}_i}{\hat{\mathbf{b}}_i}\right) - \log g_\phi\left(\frac{\mathbf{x}_i - \hat{\boldsymbol{\mu}}_i}{\hat{\mathbf{b}}_i}\right) + \log \hat{\mathbf{b}}_i. \quad (5)$$

For more details, please refer to the original RLE paper. We estimate the uncertainty ε for a predicted keypoint $\hat{\boldsymbol{\mu}}$ as the quantiles at a pre-set level α as

$$\int_{\hat{\boldsymbol{\mu}} - \varepsilon}^{\hat{\boldsymbol{\mu}} + \varepsilon} p_{\theta, \phi}(\mathbf{x}|\mathbf{z}) d\mathbf{x} = \alpha. \quad (6)$$

Vertebra classification We consider the classification of vertebrae as a separable product of keypoint features and visual features

$$p(c|\mathcal{I}, \mathbf{x}) = p_{kps}(c|\mathbf{x}) p_{img}(c|\mathbf{z}). \quad (7)$$

where $p_{img}(c|\mathbf{z})$ is the output of a single-layer neural network from the visual features $\mathbf{z}$, and $p_{kps}(c|\mathbf{x})$ is a probability from a GSQ-based differentiable classification. We construct a *fuzzy* decision function classifier f_{kps} by approximating the hard conditions of GSQ with differentiable counterparts. We replace operations “greater than” and “less than” with a sigmoid with a threshold t , and logical “and” with the minimum operation, as per the principles of e.g. [16], yielding soft class predictions $c|\mathbf{x} \sim p_{kps}$ given the vertebra keypoints. For training the classification model, we use a weighted cross-entropy loss

$$\mathcal{L}_c = - \sum_i \sum_c w_c \mathbf{I}_c (\lambda_{kps} \log p_{kps}(c|\mathbf{x}_i) + \lambda_{img} \log p_{img}(c|\mathbf{z}_i)), \quad (8)$$

where $\lambda_{img}, \lambda_{kps}$ are hyperparameters, w_c are class weights for grade and morphology, and $\mathbf{I}_c$ is an indicator function for belonging to class c . This formulation allows us to easily separate contributions to the classification, e.g. by setting $\lambda_{img} = 0$ during training, and letting $p(c|\mathcal{I}, \mathbf{x}) = p_{kps}(c|\mathbf{x})$ during inference.

3 Experiments and Results

3.1 Dataset

The Sahlgrenska University Hospital Prospective Evaluation of Risk of Bone Fractures (SUPERB) is a population-based study of 3028 older women aged 77.8 ± 1.6 from Gothenburg, Sweden [10,15]. At baseline, 13 thoracic and lumbar vertebrae were annotated with keypoints and analyzed using lateral low-dose radiographs, and prevalent fracture grade and morphology according to GSQ by two annotators with a high degree of inter-reader agreement. Due to the noisy quality of the low-dose X-ray and the expensive nature of annotation, the number of labelled vertebrae is limited, ranging from 1 to 13 per patient. Out of the full set of patients, 2919 had digitized annotations of 11,605 vertebrae. See suppl. Table 1 for an overview of the data.

Classification	$\mathcal{L}_{rle}$	$\mathcal{L}_{kps}$	$\mathcal{L}_{img}$	AUC	F_1	sensitivity	specificity
<i>normal</i> <i>vs.</i> <i>rest</i>	✓			$0.96 \pm 0.01^*$	$0.51 \pm 0.03^*$	0.90 ± 0.03	$0.90 \pm 0.01^*$
	✓	✓		$0.92 \pm 0.01^*$	$0.90 \pm 0.02^*$	$0.81 \pm 0.04^*$	0.90 ± 0.06
	✓		✓	0.98 ± 0.01	0.70 ± 0.10	$0.90 \pm 0.04^*$	0.96 ± 0.02
	✓	✓	✓	0.98 ± 0.01	0.67 ± 0.07	0.92 ± 0.04	0.95 ± 0.02
<i>normal+mild</i> <i>vs.</i> <i>moderate+severe</i>	✓			$0.97 \pm 0.01^*$	$0.42 \pm 0.05^*$	0.93 ± 0.03	$0.90 \pm 0.03^*$
	✓	✓		$0.97 \pm 0.01^*$	$0.45 \pm 0.08^*$	0.91 ± 0.03	$0.91 \pm 0.03^*$
	✓		✓	0.98 ± 0.01	0.71 ± 0.09	0.928 ± 0.032	0.97 ± 0.01
	✓	✓	✓	0.99 ± 0.01	0.73 ± 0.08	0.934 ± 0.039	0.97 ± 0.01

Table 1: *Ablation study of loss components.* 5-fold comparison of the vertebra grade classification on the test set ($N_{vertebra} = 1743$). F_1 , specificity and sensitivity computed at the Youden operating point. (*) indicates significant difference to the full model using a Bonferroni t -test at $\alpha = 0.05$.

3.2 Model implementation and training

Models were implemented in PyTorch 2.2 on two NVIDIA RTX A6000 GPUs. All metrics were computed using the `scikit-learn` library, v. 1.3.0. We randomly held out 438 out of 2919 patients (15%) for final testing, and divided the remaining 2481 (85%) patients into 5 folds for cross-validation for model training and tuning. Due to the low presence of crush deformities (suppl. Table 1), wedge and crush classes were combined into a *wedge-like* class.

Missing labels (vertebrae) is often detrimental for detection models [23]. Here we imputed missing keypoints using k -nearest neighbours before training, and downweighed their contribution to eq. (1) by a factor 10^{-3} . Models were trained with the Adam optimizer. DINO was first trained using the full data with imputations, and then fine-tuned with just the fully annotated data. Full images were downsampled to half the original size and padded to a uniform size per batch, and vertebra bounding boxes were resized to (224, 224). Standard augmentation techniques such as rotations, flipping, and jittering were applied. For complete hyperparameters chosen with grid search, see the suppl. Table 2.

3.3 Ablation study

We investigated the contributions of four combinations to eq. (8): keypoint regression ($\lambda_{img} = \lambda_{kps} = 0$), GSQ-guided loss ($\lambda_{img} = 0, \lambda_{kps} = 1$), classification from image features ($\lambda_{img} = 1, \lambda_{kps} = 0$) and combining both ($\lambda_{img} = 1, \lambda_{kps} = 1$). Results from the ablation study are shown in Table 1 calculated as mean and standard deviation on the test set over the 5 models from each fold. AUC, F_1 , sensitivity and specificity are computed as one-vs.-rest binary classification for *normal* vertebrae vs. fractures and *normal+mild* vs. *moderate+severe*. See the supplementary for results on morphology and confusion matrices.

Our proposed method, using both the rule-based loss and image feature loss outperforms the rest on identifying the clinically most relevant grouping of *normal+mild* vertebrae, with an AUC of 0.99. While using the image-based loss

Method	Level	Ref.	Expl.	AUC	F_1	sensitivity	specificity
DINO	vert.	-	X	0.95 ± 0.02	$0.33 \pm 0.06^*$	0.90 ± 0.05	$0.85 \pm 0.06^*$
CNN	vert.	[6]	X	0.97 ± 0.01	0.72 ± 0.04	0.90 ± 0.04	0.976 ± 0.004
RF	vert.	[3]	X	0.95 ± 0.02	0.56 ± 0.11	0.85 ± 0.03	0.95 ± 0.03
XVFA (<i>ours</i>)	vert.	-	✓	0.97 ± 0.02	0.65 ± 0.11	0.90 ± 0.06	0.96 ± 0.02
CNN	pat.	[18,17]	X	0.81 ± 0.09	0.62 ± 0.11	$0.60 \pm 0.16^*$	0.90 ± 0.04
XVFA (<i>ours</i>)	pat.	-	✓	-	0.70 ± 0.10	0.92 ± 0.05	0.86 ± 0.08

Table 2: *SOTA comparison.* 5-fold comparison of SOTA methods on test set on vertebra detection and classification (*vert.*, $N_{vertebra} = 1743$) and patient/image-level classification (*pat.*, $N_{patient} = 438$) for *normal* vertebrae+*mild* fractures vs. *moderate*+*severe* fractures. (*) indicates significant difference to our model using a Bonferroni *t*-test at $\alpha = 0.05$.

yields competitive results on its own, it retains the issues of explainability and lacks connection to GSQ. Fracture predictions may be visualized in terms of decision borders for MPR and MAR, see Figure 3b-c, showing that the classification of individual vertebrae in the test set closely align with GSQ, providing a useful tool to find disagreements between image and keypoint-based classifications.

3.4 SOTA model comparison

In the end-to-end results, we compared results on vertebra level and on patient-aggregated level using our best performing model ablation. We implemented a similar CNN model to [18,17], using a pre-trained Inception-ResNet-v2, classifying max severity grade in the image (learning rate of 10^{-5} , batch size 16), comparing with XVFA by aggregating per patient post-classification. As a baseline for vertebra-level detection, we trained DINO models with the same settings as earlier, but classifying grade. We compared with [6] (equiv. to our method with $\lambda_{kps} = 0$) and [3], using a random forest (100 trees, depth 2) on vertebra features h_a, h_m, h_p , APR, MPR, MAR. The results are shown in Table 2.

XVFA outperforms the SOTA and DINO baseline on the important screening metrics, AUC and sensitivity, and achieves comparable results on the rest. Like in the ablation study, using a convolutional image classifier is competitive, but still lacks the explicit relation to GSQ. It is notable that the approach of [18,17] vastly underperforms on our images, likely due to the low signal-to-noise ratio of the low-dose radiograph images. Lastly, as demonstrated by this paper, decision trees can solve the fracture classification task well, and it is no surprise that the random forest approach has competitive results. However, as is also the case with purely image-based classification, random forests are not explainable and directly relatable to clinical methodology.




Fig. 3: *Visualizations.* (a) Test sample model prediction likelihoods (*red*) and ground truth keypoints (*green*). GSQ decision borders for (b) morphology and (c) grade, for visual explanation of model results, each point corresponding to a predicted fractured vertebra in the test set.

3.5 Uncertainty quantification

We quantified the uncertainty of individual models using the full likelihood from the normalizing flow. For the best model on the validation set, we sample 1000 keypoints with replacement from the densities, see e.g. Figure 3a acquiring a $\text{mean} \pm \text{standard deviation}$ for the model on the test set: AUC 0.985 ± 0.002 , sensitivity 0.941 ± 0.004 , specificity 0.988 ± 0.004 , and F_1 of 0.842 ± 0.029 . Lastly, we tested model alignment with human annotators by asking the annotator to re-analyze a random subset of 49 scans from the test set. The annotator had a mean intra-reader deviation of approximately ± 4.5 pixels, and our model had an uncertainty of ± 5.3 pixels within the 5th and 95th percentiles, indicating model reliability on par with human observers.

4 Discussion and Conclusion

We have shown that XVFA can accurately classify vertebral fractures on par with or better than SOTA, with a per-vertebra AUC of 99% and sensitivity of 93% and end-to-end AUC of 97%, without sacrificing explainability. Proper uncertainty quantification yields reliability measures both in keypoint localization and fracture classification. We also demonstrate that use of black box neural classifiers is not necessarily superior to using an explainable alternative based

on GSQ. Our method is also flexible, and may be updated with any SOTA detector. Moreover, while this paper has mostly focused on the established task of fracture grade classification, excellent results on morphology classification are available in the supplementary which lack comparative baselines.

However, our model does have some limitations. It does not account for additional features, such as adjacent vertebrae into classification (see e.g. [3]) and has no native means of handling missing vertebrae in annotation. A more explicit modelling of spinal structure would likely address these issues. Moreover, our approach does not distinguish *epistemic* and *aleatoric* uncertainty, focusing only on the latter. Normalizing flows may also have an issue with overfitting [11], which we did not notice in this paper, requiring research on other cohorts to validate. In spite of these shortcomings, our model displays sufficient sensitivity for screening ($> 90\%$) of fracture grades, and may provide trustworthy and explainable decision support useful in clinical work that can aid the evaluation of VFA, alleviate physician workload, and ultimately improve patient outcome.

5 Data Use Declaration and Acknowledgment

The collection of data for the SUPERB study was approved by the Ethical Review Authority Swedish Ethical Review Authority (DNR 929-12). The data is not publicly available as it contains sensitive health information. The source code and trained models can be downloaded for academic or non-commercial purposes.³ **Acknowledgments.** This study was funded in part through the AIDA project grant DNR 2021-01420. **Disclosure of Interests.** Outside of this research, ML has received lecture fees from Amgen, Astellas, Lilly, Meda, Renapharma, UCB Pharma, and consulting fees from Amgen, Radius Health, UCB Pharma, Parexel International, Renapharma and Consilient Health, LJ has received lecture fees from UCB Pharma. All other authors declare no competing interests.

³ <https://github.com/waahlstrand/xvfa>

References

1. Bürgin, V., Prevost, R., Stollenga, M.F.: Robust vertebra identification using simultaneous node and edge predicting graph neural networks. In: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*. pp. 483–493. Springer Nature Switzerland (2023)
2. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. *Lecture Notes in Computer Science* **12346 LNCS**, 213–229 (5 2020). https://doi.org/10.1007/978-3-030-58452-8_13
3. Cheng, L.W., Chou, H.H., Cai, Y.X., Huang, K.Y., Hsieh, C.C., Chu, P.L., Cheng, I.S., Hsieh, S.Y.: Automated detection of vertebral fractures from x-ray images: A novel machine learning model and survey of the field. *Neurocomputing* **566**, 126946 (1 2024). <https://doi.org/10.1016/J.NEUCOM.2023.126946>
4. Dinh, L., Sohl-Dickstein, J., Bengio, S.: Density estimation using real nvp. *arXiv preprint arXiv:1605.08803* (2016)
5. Dong, Q., Luo, G., Lane, N.E., Lui, L.Y., Marshall, L.M., Johnston, S.K., Dabous, H., O'Reilly, M., Linnau, K.F., Perry, J., Chang, B.C., Renslo, J., Haynor, D., Jarvik, J.G., Cross, N.M.: Generalizability of deep learning classification of spinal osteoporotic compression fractures on radiographs using an adaptation of the modified-2 algorithm-based qualitative criteria. *Academic radiology* **30**, 2973–2987 (12 2023). <https://doi.org/10.1016/J.ACRA.2023.04.023>
6. Dong, Q., Luo, G., Lane, N.E., Lui, L.Y., Marshall, L.M., Kado, D.M., Cawthon, P., Perry, J., Johnston, S.K., Haynor, D., Jarvik, J.G., Cross, N.M.: Deep learning classification of spinal osteoporotic compression fractures on radiographs using an adaptation of the genant semiquantitative criteria. *Academic Radiology* **29**, 1819–1832 (12 2022). <https://doi.org/10.1016/j.acra.2022.02.020>
7. Genant, H.K., Wu, C.Y., van Kuijk, C., Nevitt, M.C.: Vertebral fracture assessment using a semiquantitative technique. *Journal of Bone and Mineral Research* **8**, 1137–1148 (9 1993). <https://doi.org/10.1002/JBMR.5650080915>
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* **2016**, 770–778 (12 2015). <https://doi.org/10.1109/CVPR.2016.90>
9. Software and systems engineering - Software testing, Part 11: Guidelines on the testing of AI-based systems. Standard, International Organization for Standardization (2020)
10. Johansson, L., Sundh, D., Magnusson, P., Rukmangatharajan, K., Mellström, D., Nilsson, A.G., Lorentzon, M.: Grade 1 vertebral fractures identified by densitometric lateral spine imaging predict incident major osteoporotic fracture independently of clinical risk factors and bone mineral density in older women. *Journal of Bone and Mineral Research* **35**, 1942–1951 (10 2020). <https://doi.org/10.1002/JBMR.4108>
11. Kirichenko, P., Izmailov, P., Wilson, A.G.: Why normalizing flows fail to detect out-of-distribution data. *Advances in neural information processing systems* **33**, 20578–20589 (2020)
12. Kouchaki, S., Ding, X.R., Sanei, S., Chaddad, A., Peng, J., Xu, J., Bouridane, A.: Survey of explainable ai techniques in healthcare. *Sensors* 2023, Vol. 23, Page 634 **23**, 634 (1 2023). <https://doi.org/10.3390/S23020634>

13. Li, J., Bian, S., Zeng, A., Wang, C., Pang, B., Liu, W., Lu, C.: Human pose regression with residual log-likelihood estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 11025–11034 (10 2021)
14. Lindsay, R., Silverman, S.L., Cooper, C., Hanley, D.A., Barton, I., Broy, S.B., Licata, A., Benhamou, L., Geusens, P., Flowers, K., Stracke, H., Seeman, E.: Risk of New Vertebral Fracture in the Year Following a Fracture. *JAMA* **285**(3), 320–323 (01 2001). <https://doi.org/10.1001/jama.285.3.320>, <https://doi.org/10.1001/jama.285.3.320>
15. Lorentzon, M., Nilsson, A.G., Johansson, H., Kanis, J.A., Mellström, D., Sundh, D.: Extensive undertreatment of osteoporosis in older swedish women. *Osteoporosis International* **30**, 1297–1305 (6 2019), <https://link.springer.com/article/10.1007/s00198-019-04872-4>
16. Mamdani, E.H.: Application of fuzzy algorithms for control of simple dynamic plant. *Proceedings of the Institution of Electrical Engineers* **121**, 1585–1588 (1974). <https://doi.org/10.1049/PIEE.1974.0328/CITE/REFWORKS>
17. Monchka, B.A., Schousboe, J.T., Davidson, M.J., Kimelman, D., Hans, D., Raina, P., Leslie, W.D.: Development of a manufacturer-independent convolutional neural network for the automated identification of vertebral compression fractures in vertebral fracture assessment images using active learning. *Bone* **161**, 116427 (8 2022). <https://doi.org/10.1016/J.BONE.2022.116427>
18. Murata, K., Endo, K., Aihara, T., Suzuki, H., Sawaji, Y., Matsuoka, Y., Nishimura, H., Takamatsu, T., Konishi, T., Maekawa, A., et al.: Artificial intelligence for the detection of vertebral fractures on plain spinal radiography. *Scientific Reports* **10**(1), 20031 (2020)
19. Rezaatofghi, H., Tsoi, N., Gwak, J., Sadeghian, A., Reid, I., Savarese, S.: Generalized intersection over union: A metric and a loss for bounding box regression. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* **2019-June**, 658–666 (2 2019). <https://doi.org/10.1109/CVPR.2019.00075>
20. Rezende, D., Mohamed, S.: Variational inference with normalizing flows. In: Bach, F., Blei, D. (eds.) *Proceedings of the 32nd International Conference on Machine Learning*. *Proceedings of Machine Learning Research*, vol. 37, pp. 1530–1538. PMLR (7 2015), <https://proceedings.mlr.press/v37/rezende15.html>
21. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. *International Journal of Computer Vision* **128**, 336–359 (10 2016). <https://doi.org/10.1007/s11263-019-01228-7>
22. Wu, H., Zhang, J., Fang, Y., Liu, Z., Wang, N., Cui, Z., Shen, D.: Multi-view vertebra localization and identification from ct images. In: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*. pp. 136–145. Springer Nature Switzerland, Cham (2023)
23. Xu, M., Bai, Y., Ghanem, B., Liu, B., Gao, Y., Guo, N., Ye, X., Wan, F., You, H., Fan, D., et al.: Missing labels in object detection. In: *CVPR workshops*. vol. 3 (2019)
24. Zhang, H., Li, F., Liu, S., Zhang, L., Su, H., Zhu, J., Ni, L.M., Shum, H.Y.: Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605* (2022)