

# Enhancing Interpretability of Vertebrae Fracture Grading using Human-interpretable Prototypes

Poulami SINHAMAHAPATRA [HTTPS://ORCID.ORG/0000-0002-3873-9623](https://orcid.org/0000-0002-3873-9623) poulami.sinhamahapatra@iks.fraunhofer.de  
Fraunhofer IKS and Technical University of Munich, Germany

Suprosanna SHIT [HTTPS://ORCID.ORG/0000-0003-4435-7207](https://orcid.org/0000-0003-4435-7207) suprosanna.shit@tum.de  
Technical University of Munich, Germany

Anjany SEKUBOYINA [HTTPS://ORCID.ORG/0000-0002-5601-284X](https://orcid.org/0000-0002-5601-284X) anjany.sekuboyina@tum.de  
Technical University of Munich, Germany

Malek El HUSSEINI [HTTPS://ORCID.ORG/0000-0002-2952-3708](https://orcid.org/0000-0002-2952-3708) malek.husseini@tum.de  
Technical University of Munich, Germany

David SCHINZ [HTTPS://ORCID.ORG/0000-0003-3734-1135](https://orcid.org/0000-0003-3734-1135) david.schinz@tum.de  
Technical University of Munich, Germany

Nicolas LENHART nicolas.lenhart@tum.de  
Technical University of Munich, Germany

Bjoern MENZE [HTTPS://ORCID.ORG/0000-0003-4136-5690](https://orcid.org/0000-0003-4136-5690) bjoern.menze@uzh.ch  
University of Zurich, Switzerland

Jan KIRSCHKE [HTTPS://ORCID.ORG/0000-0002-7557-0003](https://orcid.org/0000-0002-7557-0003) jan.kirschke@tum.de  
Technical University of Munich, Germany

Karsten ROSCHER [HTTPS://ORCID.ORG/0000-0002-9458-104X](https://orcid.org/0000-0002-9458-104X) karsten.roscher@iks.fraunhofer.de  
Fraunhofer IKS, Munich, Germany

Stephan GUENNEMANN [HTTPS://ORCID.ORG/0000-0001-7772-5059](https://orcid.org/0000-0001-7772-5059) guennemann@in.tum.de  
Technical University of Munich, Germany

## Abstract

Vertebral fracture grading classifies the severity of vertebral fractures, which is a challenging task in medical imaging and has recently attracted Deep Learning (DL) models. Only a few works attempted to make such models human-interpretable despite the need for transparency and trustworthiness in critical use cases like DL-assisted medical diagnosis. Moreover, such models either rely on post-hoc methods or additional annotations. In this work, we propose a novel interpretable-by-design method, ProtoVerse, to find relevant sub-parts of vertebral fractures (prototypes) that reliably explain the model's decision in a human-understandable way. Specifically, we introduce a novel diversity-promoting loss to mitigate prototype repetitions in small datasets with intricate semantics. We have experimented with the VerSe'19 dataset and outperformed the existing prototype-based method. Further, our model provides superior interpretability against the post-hoc method. Importantly, expert radiologists validated the visual interpretability of our results, showing clinical applicability.

**Keywords:** Machine Learning, Interpretability, Explainability




Figure 1: Prototype-activated regions (yellow boxes) on the fractured vertebra of a typical test sample provides human-interpretable reasoning for its fracture grade G3.

## 1. Introduction

Osteoporosis is one of the most prevalent diseases among elderly people, which is clinically manifested as bone fractures localised typically in the spine, hip, distal forearm, and proximal humerus. These fractures can lead to severe pain, kyphosis, disability, as well as approx. 12-fold increased risk for further fractures and 8-fold increased mortality (Cauley et al., 2000). If such fractures are detected timely, they can be effectively treated to reduce further fracture occurrences and morbidity. Unfortunately, most such cases go undetected until the occurrence of the first acute symptoms. Vertebral Compression Fractures (VCFs) are the most prevalent among osteoporotic fractures occurring in the vertebrae. VCFs inflict 30 – 50% population above the age of 50 (Ballane et al., 2017) and have been reported to be undiagnosed in 84% cases via Computed Tomography (CT) (Carberry et al., 2013). In clinical practice, radiologists primarily rely on inspecting the sagittal plane of CT scans for detecting VCFs. Moreover, radiologists grade the detected VCF using *Genant scale* (Genant et al., 1993), which measures the reduction in vertebral height from CT images. The scale has four grades: G0 (Healthy), G1 (Mild/ < 25% height reduction), G2 (Medium/ 25 – 40% reduction), and G3 (Severe/> 40% reduction). Note that, fracture grading is a more complex task than fracture detection since it is a fine-grained classification and relies on radiologists’ skill and experience. Alternatively, Deep Learning (DL) has the potential to be used as an assistive tool for overloaded radiologists to accelerate VCF detection and VCF grading. However, the conventional DL-based method offers slim to none interpretability, making its deployment unreliable in safety-critical systems such as clinical practice. In order to effectively reduce the burden of radiologists and facilitate faster diagnosis, human-interpretable DL methods are the need of the hour. In this paper, we aim to develop a novel interpretable-by-design DL model for explainable VCF grading. Recently eXplainable AI (XAI) methods have been developed to estimate decision-making factors of Deep Neural Networks (DNNs) into human-interpretable formats from the learnt representations or activations. However, what makes an explanation sufficiently human-interpretable is largely subjective of the XAI method, use-case, and the dataset in question (Sinhamahapatra. et al., 2023). While some initial approaches rely on post-hoc explanations, recent methods are designed to be inherently interpretable. For critical healthcare applications, it is important that we generate explanations that are aligned with how doctors would interpret them. A growing direction in current XAI approaches is to learn representations that can be tied directly to higher level *human-understandable concepts*. An example from natural images

could be predicting an image of a *red-billed hornbill*, which depends on the presence of concepts like *red bill*. However, an explicit identification of such concepts requires costly expert annotations for each sample and often can be non-exhaustive. Instead, one could learn such concepts in an unsupervised manner as *prototype* parts, which are semantically recurring features for a given class in a dataset. These prototypes are representative parts sufficient to identify the given class, e.g., if prototype parts of *beaks*, *wings*, *tails* are learnt, they could together reliably predict a *bird*. Similarly, in VCF fracture grading (c.f. Fig. 1), detecting G3 grade fracture relies on severe deformities in the vertebral body. While learning such prototype representations in an unsupervised manner is a challenge by itself, visualising such implicitly learnt representations in a human-understandable way is also difficult. Recently *Prototypical Parts Network* (ProtoPNet) (Chen et al., 2019), has been introduced that enables DNNs to learn internal representations as a set of interpretable prototypes consisting of parts of the closest training image. During inference, parts of test samples could be visually reasoned to be similar (‘this-looks-like-that’) to some of its closest prototypes. In this work, we leverage prototype-based learning for human-interpretable VCF grading. By learning visual prototypes that sufficiently represent different fracture grades, one can explain the inferred VCF grading of a test sample using its similarity with the prototypes from that fracture grade (c.f. Fig. 1). Such an explainable design would bring trustworthiness to the deployed system and reduce the workload of clinicians and diagnosis time for patients. However, existing prototype-based models struggle to capture necessary visual and semantic variations within the prototypes of the same class. Repetition in prototypes can result in inadequate human-interpretable visual explanations during inference. Crucially, these methods were designed to target natural images where the semanticity is intuitive and abundant. In comparison, semanticity in medical images is intricate and involves a complex, expert understanding of the normality of anatomical patterns, including edges, shapes, volume, and local contrast. Hence, these models are more prone to visual feature collapse in the prototype layer for the medical dataset, which calls for appropriate customisation. This is even exaggerated in the case of small dataset size. We mitigate this by a novel diversity-promoting loss, which forces the network to focus on a diverse set of vertebral features, e.g., various fracture regions. Further, we employ imbalance-aware classification loss to alleviate dataset imbalance. In the end, our proposed method is suitably applied for VCF grading in an imbalanced and small dataset.

**Contributions:** In summary, our contributions are:

- We propose a novel interpretable-by-design method *ProtoVerse*, for VCF grading based on prototype parts learning. To the best of our knowledge, we are the first to leverage interpretability based on prototypes into clinical VCF datasets.
- We propose a novel objective function to enhance the visual and semantic diversity of the learnt prototype for each fracture / healthy class. Additionally, we mitigate huge class imbalances with the help of differential weighting in the classification loss.
- Our method substantially outperforms ProtoPNet, a state-of-the-art prototype-based method. Moreover, our method precisely localises clinically relevant parts in the vertebral body, while a prominent post-hoc method, GradCAM, fails to achieve this.

- Finally, we demonstrate the clinical applicability of our method by validating its human-interpretability with the help of expert radiologists.

## 2. Related Work

**Human-interpretable DL Models:** A collection of XAI works for visual tasks can be found in the surveys (Nguyen et al., 2019; Samek and Müller, 2019) including the medical application (Tjoa and Guan, 2020). The explanations can be categorised based on several criteria, such as (a) local vs. global, (b) post-hoc vs. interpretable-by-design (IBD), and (c) explicitly specified vs. implicitly derived. Depending upon the user requirement, a particular XAI approach can be used for generating explanations. *Local* explanations deal with single sample by looking at its specific parts with highest attribution for a particular decision in DNNs, such as saliency-based class activation map (CAM) (Zhou et al., 2016), GradCAM (Selvaraju et al., 2017) and its variants like GradCAM++ (Chattopadhyay et al., 2017) and XGradCAM (Fu et al., 2020), gradient-based local sensitivity, (Sundararajan et al., 2017) and relevance back-propagation (Bach et al., 2015). Previous studies have pointed out that saliency-based methods for post-hoc explainability are limited under constant intensity shift (Kindermans et al., 2019), model weight randomization (Arun et al., 2021), relatively smaller regions of interest (Saporta et al., 2022), confounders (Sun et al., 2023), and often disagree with medically correct explanation (Rafferty et al., 2022). In contrast, *global* explanations consider the whole dataset and generate explanations based on recurring features in multiple samples. Global explanations derived from such learnt internal representations are more useful and often necessary for tasks like anomaly detection. Most local methods generate *post-hoc* explanations from already trained models. Since these explanations are generated a posteriori, they are sub-optimal to make use of high-level semantic recurrence in the internal representation. Therefore, in order to make the internal representations human-interpretable, we need to consciously design global methods that are inherently *IBD* (Rudin, 2019). Here, one recent approach maps human-understandable concepts to internal representations as an interpretable bottleneck layer (Schwalbe, 2022). These concepts can be either *explicitly specified* per sample as attributes in natural language by domain experts as in concept learning models (Koh et al., 2020; Kazhdan et al.; Fong and Vedaldi, 2018; Fang et al., 2020; Cao et al., 2021) or *implicitly derived* as in prototype-based learning models (Chen et al., 2019; Nauta et al., 2021; Gautam et al., 2021; Li et al., 2018; Donnelly et al., 2022; Rymarczyk et al., 2022). The latter does not depend on concept-specific annotations but rather simply learns to associate latent representations to parts of the closest training images. These parts are termed as *prototypes* (Chen et al., 2019), which serve as semantically relevant visual explanations. Subsequent works have pointed out multiple limitations, such as the semantic gap between latent and input space (Hoffmann et al., 2021), and the distance of prototypes from the decision boundary (Wang et al., 2023). Since then, many works have extended and improved prototypical learning by imposing consistency and stability criteria (Huang et al., 2023), using support prototypes (Wang et al., 2023), mining negative prototypes (Singh and Yow, 2021), and receptive field-based localization (Carmichael et al., 2024).

**Vertebral Fracture Detection:** Numerous recent works have studied vertebrae fracture detection on both the patient and vertebrae levels; however, only a few have attempted

to make the network decision IBD. The CNN-based networks take center stage in the architectural choice, which is used in both 2D slices and 3D volumes of CT images. Among 2D-based works, a custom contrastive loss based on grading hierarchy to process the reformatted mid-sagittal 2D slices was proposed in (Husseini et al., 2020). In (Pisov et al., 2020), a key points based method was proposed to localize vertebrae and subsequently estimate fracture severity. An LSTM-based approach was used in (Tomita et al., 2018) and (Bar et al., 2017) to leverage inter-slice and intra-slice feature dependencies, respectively. While 2D methods can efficiently capture discriminative features of vertebrae, 3D methods explicitly target 3D context at the cost of increased computational cost. A 3D dual-pathway architecture for dense voxel-wise fracture detection was proposed in (Nicolaes et al., 2020). A sequence-to-sequence CNN for inter-patch feature accumulation was used in (Chettrit et al., 2020). A patient-level feature aggregation from 3D patches was proposed in (Yilmaz et al., 2021). A two-stream architecture to explore discriminative features for fine-grained classification was used in (Feng et al., 2021).

Recently, (Engstler et al., 2022b) proposed to leverage an activation map to find correlated image regions to a fracture grade by expert validation. However, their method offers post-hoc interpretability and heavily depends on the post-processing heuristics. In stark contrast, we aim to obtain a global IBD model which enforces the network to learn the most semantically recurring features across the dataset. (Zakharov et al., 2023) proposes a fracture grading pipeline, which provides interpretability based on the measured heights between landmarks of vertebrae. While this approach closely follows the Genant definition, it requires additional annotation of six key points per vertebrae, making the annotation process complex and more time-consuming.

Few recent approaches applied prototype-based learning methods in medical use-cases, such as Alzheimer’s disease (Mohammadjafari et al., 2021), brain tumor classification (Wei et al., 2023), and breast masses classification (Carloni et al., 2022). However, their goal is limited to binary classification, such as anomaly detection (healthy vs. unhealthy). In comparison, our method brings interpretability to fine-grained multi-class medical classification tasks such as VCF grading, where small sample sizes and severe class imbalances pose additional challenges. Previously, (Kim et al., 2021) has shown multi-class use cases on chest X-rays; however, their approach is tailored for disease with dynamic areas of focus.

### 3. Dataset

In this work, we use the VerSe’19 dataset from the VerSe 2019 challenge for VCF grading. VerSe (Sekuboyina et al., 2021) is a large-scale, multi-detector, multi-site CT spine dataset comprising 374 scans from 355 patients. To minimize existing inter-rater variability of publically available annotation (Löffler et al., 2020), we updated annotations in-house by a group of radiologists with 15 years of median experience. For VCF, in this work, we consider healthy, G2 and G3. Following (Husseini et al., 2020; Keicher et al., 2023; Engstler et al., 2022a), we do not consider G1 since it is often clinically challenging to differentiate G1 from healthy class and is prone to significant inter-rater variability. We have altogether 1444 vertebrae annotation; out of 1308 are healthy, 76 are G2, and 52 are G3. We have 1086 training samples and 358 test samples. We have 326 healthy, 18 G2, and 14 G3 test

samples. Clearly, our main challenge is to devise a useful learning strategy for a successful interpretable DL method to mitigate high class imbalance and small sample size.

#### 4. Methodology: ProtoVerse

In this section, we propose a novel interpretable method called ProtoVerse for vertebral fracture grading. At the core of this lies an interpretable architecture motivated by ProtoPNet (Chen et al., 2019), which is an image classifier network that learns inherent representations for relevant sub-parts of an image. Once these representations are learnt, they serve as visual explanations for the model output. To obtain such visual explanation, sub-parts or patches of training images closest to the learnt representations are stored in a set of visual patches called *prototypes* to enhance human-interpretation of the inner workings of DNN models.

Figure 2 shows the overview of the *ProtoVerse* for learning visual prototypes for vertebral fractures. The model consists of a convolutional feature extractor backbone  $f$  followed by two additional  $1 \times 1$  layers, a prototype layer  $g_p$ , and a fully connected layer  $h$  with weights  $w_h$ . The feature extractor is usually taken from standard Convolutional Neural Networks (CNNs), e.g., ResNets, DenseNets, and VGG. These CNNs are initialised using ImageNet pre-trained weights. During training, firstly, we learn the prototype vectors in the latent space via stochastic gradient descent over the objective function until the last layer. Next, the last layer for classification is updated by a convex optimisation. Finally, the learnt prototype vectors are visualised through retrieved *prototype* patches/sub-parts of training images. In the following, we expand upon each of the working stages of ProtoVerse.

##### 4.1 Learning Prototype Vectors

Let us consider an input image  $\mathbf{x}$  with a fracture grade  $\in \{G0, G2, G3\}$ . The feature extractor generates a convolutional output  $\mathbf{z} = f(\mathbf{x})$  with a shape of  $\mathbf{z} \in \mathbb{R}^{H \times W \times D}$ , where  $D$  is the feature size in the last convolutional layer. Given the feature extractor is ImageNet pre-trained, the input image is resized to  $224 \times 224 \times 3$ . The output of the  $f$  produces  $7 \times 7$  patches, i.e.  $H = 7, W = 7$ .

Next, the output of  $f$  goes through the prototype layer  $g_p$  consisting of learnable prototype vectors  $\mathbb{P} = \{\mathbf{p}_j^k\}_{j=1:c, k=1:m}$ , where  $m$  is the number of prototypes per class and  $c$  is the total number of classes. Thus, the total number of prototypes  $n = m \times c$ . Each prototype  $\mathbf{p}_j^k \in \mathbb{R}^{h \times w \times D}$  has the same feature size as the last convolutional layer. Following ProtoPNet (Chen et al., 2019), we have taken prototypes with  $h = w = 1$ . Thus, each prototype corresponds to a prototypical activation pattern  $1 \times 1 \times D$  patch in the  $7 \times 7 \times D$  last-layer feature. Since the  $7 \times 7 \times D$  feature vector represents the entire input image, a prototype vector  $\mathbf{p}_j^k$  corresponds to a proportionate field of view of a single feature in a patch, which is considered as prototype patch/sub-part of the input image. Thus,  $n$  number of prototype vectors corresponds to the  $n$  most salient image parts characterising the given dataset.

Next, we calculate the  $\ell_2$  distances between each of the prototype  $\mathbf{p}_j^k$  and all the similar-sized patches of the feature output layer  $\mathbf{z}$ . The smaller the distance, the higher the similarity of the prototype vector to the feature patch for the given input image. Thus, the point-wise inversion of the  $\ell_2$  distances provides  $7 \times 7$  activation map referred to as *similarity*

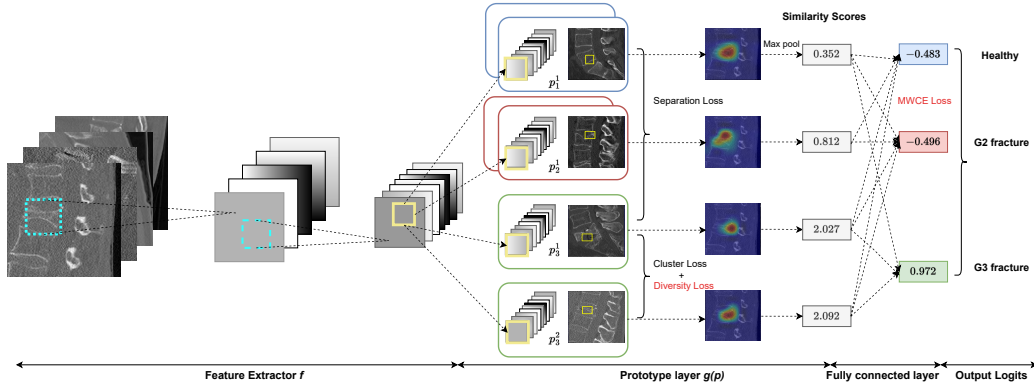


Figure 2: ProtoVerse architecture for learning prototypes for VCF grading. Prototypes from each class are shown as : healthy (blue), G2 (red), and G3 (green), which learns representative image parts for each class through separation and clustering loss. For example, healthy prototypes emphasise straight vertebral edges, while G2 and G3 prototypes capture the different degrees of deformities in vertebrae. Notably, our novel Diversity loss ensures the capture of visual variations within a class such as  $p_3^1$  and  $p_3^2$  highlighting different fracture regions in G3. Given a G3 input test sample, the prototype patch  $p_3^1$  and  $p_3^2$  belonging to G3 shows the strongest presence (similarity score 2.027 and 2.092) in various fracture regions for vertebrae of interest. While the prototype parts  $p_2^1, p_1^1$  belonging to G2 and healthy classes have lesser similarity score. Final classification logits are trained with our MWCE loss to mitigate class imbalance.

*heatmap*. To reduce the similarity heatmap into a single *similarity score*, global max pooling is performed. This helps in interpreting how strongly a particular prototype vector matches with a specific part of the input image. Thus, the prototypes for a particular class  $j$  are responsible for capturing the most relevant parts for images of class  $j$ . Lastly, in the fully connected layer  $h$ , the  $m$  similarity scores from each class are multiplied with the weight matrix  $w_h$ , which constitutes the *class connection score*. For an input  $x_i$ , the output logits are obtained as  $\hat{y}_i = h \circ g_p \circ f(x_i)$ . The resultant outputs are normalised using softmax to produce the probabilities for images belonging to various classes.

To learn meaningful prototype vectors in the latent space, prototypes of each class should be clustered together and well-separated from the prototypes of different classes. ProtoPNet (Chen et al., 2019) tried to achieve this goal by introducing intra-class clustering and inter-class separation loss, in addition to classification loss. For a given training dataset,  $\mathbb{D}_{\text{train}} = \{\mathbf{x}_i, y_i\}_{i=1}^d$  and a set of prototype vectors for  $j^{\text{th}}$  class  $\mathbb{P}_{y_i} = \{\mathbf{p}_{y_i}^k\}_{k=1}^m$ , the cluster loss ( $\mathcal{L}_{\text{clst}}$ ) and separation loss ( $\mathcal{L}_{\text{sep}}$ ) is as follows:

$$\mathcal{L}_{\text{clst}} = \frac{1}{d} \sum_{i=1}^d \min_{j: \mathbf{p}_j^k \in \mathbb{P}_{y_i}} \min_{z \in \chi(z_i)} \|z - \mathbf{p}_j^k\|_2^2; \quad \mathcal{L}_{\text{sep}} = \frac{1}{d} \sum_{i=1}^d \min_{j: \mathbf{p}_j^k \notin \mathbb{P}_{y_i}} \min_{z \in \chi(z_i)} \|z - \mathbf{p}_j^k\|_2^2 \quad (1)$$

where  $\chi(z_i)$  denotes all the patches from feature  $z_i$ . Separation loss pushes apart prototypes of different classes in the prototype layer. On the contrary, cluster loss tries to increase semantically similarity among prototype patches of its own class based on  $\ell_2$  distances. At

the same time, the weights of the fully connected layer are initialised following (Chen et al., 2019). Specifically for a given class, we initialise class connection for the prototype of the same class as 1.0 and  $-0.5$  for prototypes of a different class.

Since VerSe’19 is a relatively small dataset, it is increasingly challenging to learn relevant prototypes for fracture classes. Particularly, cluster loss in Eq. 1 encourages prototypes within the same class to be close to each other. However, due to the absence of many diverse samples, this inevitably leads to prototypes that are semantically similar and repeating patches. For a more interpretable diagnosis, it is often preferred to look at different patches of the given fracture grade, for e.g., prototypes focusing on fractures on the top/ bottom edge or the frontal/ posterior vertical edges of the vertebrae. To this end, we propose a novel loss to promote the diversity of prototypes for a given class. Moreover, VerSe’19 has only 5% fracture samples compared to healthy samples. With such a heavy class imbalance, the classification and Separation loss tends to be skewed towards healthy class. Since interpretable learning demands a sufficient number of samples, we aim to keep all available healthy vertebrae and employ a class imbalance aware loss function. Next, we describe our proposed training strategies comprising novel prototype diversity loss and imbalance-aware classification loss for better interpretability.

#### 4.1.1 NOVEL PROTOTYPE DIVERSITY LOSS

Given very few samples from fracture classes, only a limited variation of visual parts showing predominant characteristics for the fracture classes are available to learn the prototypes. Further, semantics in medical images involve an expert understanding of anatomical patterns, including edges, shapes, and local texture. However, these features are repetitive across different scales and image locations. For example, vertebral process regions in the posterior part could have similar geometric and texture features in an anterior vertebral body. Under such constraints, clustering loss makes the network susceptible to prototype repetition. Therefore, an explicit constraint in the form of diversity loss is helpful to enforce the network to cover fine-grained variations of edges, shapes, and local textures within a certain class. Thus, we propose a novel loss component called *Diversity Loss* ( $\mathcal{L}_{\text{div}}$ ), which minimizes the absolute cosine distance between prototype vectors of the same class. In action, Diversity loss prevents the collapsing of multiple prototypes of the same class into identical ones. This is achieved by balancing the weight between Cluster and Diversity loss. As a result, the network is encouraged to find prototypes as the closest match to semantically diverse parts of the vertebrae. Experimentally, we find Diversity loss significantly mitigates the problem of redundancy (Sinhamahapatra. et al., 2023) and improves the semantic and visual diversity of the prototypes. For example, in Figure 2 the prototypes  $\mathbf{p}_3^1, \mathbf{p}_3^2$  from Grade 3 fracture focus on two distinctly different parts. It is given by the squared mean of the dot product of the prototype vectors for a given class with a coefficient ( $\lambda_{\text{div}}$ ) that decides the weight of this loss in the total loss.

$$\mathcal{L}_{\text{div}} = \frac{2}{cm(m-1)} \sum_{j=1:c} \sum_{k,l=1:m, k \neq l} (\mathbf{p}_j^k \times \mathbf{p}_j^l)^2 \quad (2)$$

where  $m$  is the number of prototypes per class, and  $c$  is the total number of classes.



#### 4.1.2 CLASS IMBALANCE AWARE CLASSIFICATION LOSS

With imbalanced dataset, the class-wise accuracy is heavily affected by model’s bias towards dominant class despite achieving high sample-wise accuracy. This leads very few samples from fractured classes to make a positive connection with the fully connected layer and in turn the prototypes learnt for these classes are not sufficiently representative for the given class. In order to provide more opportunity for learning from the samples of fractured class, we adopt differential weighting in the CE loss. However, choosing correct weighting strategy is crucial and varies with dataset size and complexity. Besides resampling strategies for tackling imbalance, we have experimented with other weighting strategies such as Inverse Number of samples (INS), Inverse of Square Root Number of Samples (ISNS). We found median weighting to be the most beneficial in our case. We have employed median-weighted cross-entropy (MWCE) in the classification loss and this has significantly improved the class-wise accuracy. Given  $d_j$  represent the number of training samples from  $j$ -th class, the weight for the  $j$ -th is calculated as follows:

$$\mathcal{L}_{\text{MWCE}}(\hat{y}_i, y_i) = - \sum_{j=1}^c w_j \cdot y_i^j \log(\hat{y}_i^j); \text{ where } w_j = \frac{\text{median}(\{d_l\}_{l=1:c})}{\sum_{l=1:c; l \neq j} d_l} \quad (3)$$

Finally, we obtain the total loss function for our ProtoVerse model as follows:

$$\mathcal{L}_{\text{Total}} = \frac{1}{d} \sum_{i=1}^d (\mathcal{L}_{\text{MWCE}} + \lambda_{\text{clst}} \mathcal{L}_{\text{clst}} + \lambda_{\text{sep}} \mathcal{L}_{\text{sep}} + \lambda_{\text{div}} \mathcal{L}_{\text{div}}) \quad (4)$$

## 4.2 Visualising Prototype Vectors

In total, we have  $n$  number of  $D$ -dimensional prototype vectors in latent space which is commonly visualised as clusters via various visualisation methods like TSNE, UMAP etc. However, to make our models interpretable, we need to visualise which part of input images are reflected into those latent prototype vectors. To address this, we follow (Chen et al., 2019) and extract the exact patch or image sub-part that activates the most in similarity score heatmap for each prototype vector. Such obtained image part level visualisations provide detailed interpretability in a human-understandable manner. For each prototype vector, the visual prototype patches are extracted based on the closest training image patch from each class based on the minimum  $\ell_2$  distances as given in Section 4.1. The above operation of finding the closest patch for all given training image patches for the class is given as below:

$$\mathbf{p}_j^k \leftarrow \arg \min_{\mathbf{z} \in \chi(\mathbf{z}_i)} \|\mathbf{z} - \mathbf{p}_j^k\|_2; \forall i \text{ s.t. } y_i = j \quad (5)$$

However, prototype patches saved are not exactly the upscaled size corresponding to  $1 \times 1 \times D$  feature patch. Instead the entire similarity heatmap for the chosen training image is upscaled to the input image dimensions. Finally, those image regions are selected as prototype patch which enclose atleast 98% in the similarity heatmap. Thus, the prototypes generated during training, act as visually interpretable parts crucial for a given class. During inference, similarity heatmap highlight the parts of a test sample that are close to the learnt prototype for a given class.

## 5. Experimental Results and Discussion

In this section, we demonstrate how prototype-based learning methods can be applied for enhanced interpretability to the vertebrae fracture grading usecase. We first provide details about dataset preparation experimental setup and then show our quantitative and qualitative results. Importantly, we provide insights from clinicians by conducting experiments to find the usefulness of such a method for the interpretability of fracture grades.

**Data Preparation:** We adopt similar data preparation steps as (Husseini et al., 2020). Since radiologists prefer sagittal views, we gather information from 3D CT volume into reformatted 2D sagittal slices. Specifically, we construct a spline along the vertebral centroids and extract 2D reformation of the vertebrae for the sagittal plane along which the spline goes through. As a result, we have a 2D reformatted slice along the mid-vertebral plane. Finally, we crop a 112x112 size region of interest comprising valuable context of at least one vertebra above and below, keeping the vertebrae of interest at the center, and subsequently resize the image to 224x224. We have not noticed any significant difference in accuracy drop in training from scratch for different image resolutions [112x112, 224x224]. However, for improved performance, we have always used an imagined-trained backbone in our experiments, and hence we have resized it to 224x224. Moreover, we provide an additional channel with a Gaussian around the centroid of the vertebra of interest.

**Experimental Setup:** We have used similar settings as in (Chen et al., 2019) for model training and prototype visualisation. In our experiments, we select a classifier without a prototype layer as a baseline for direct comparison between post-hoc interpretability of the baseline with our proposed interpretable-by-design method. Since previous methods (Husseini et al., 2020; Keicher et al., 2023) used different fracture annotations than ours, their results cannot be used for direct comparison with our model. We believe our choice of baselines with and without IBD serve as sufficient points of comparison in our experiments. Further, we show ablation of our proposed diversity loss. We have extensively evaluated with various pre-trained backbones such as VGG (Simonyan and Zisserman, 2014), ResNet (He et al., 2016), and DenseNet (Huang et al., 2017) and found that VGG-11 results in the best performance. We have also conducted an ablation study on the number of prototypes per class and found the performance decreasing with an increasing number of prototypes (Sec 5.6).

To quantitatively evaluate the performance of imbalanced dataset like VerSe’19, we have focused on *Class Average Accuracy*, *Class Average F1* scores as evaluation metrics, such that performance on individual classes are highlighted as compared to *Sample Accuracy* based on total number of correct samples.

### 5.1 Classification Results

We use 10% of the training data as a validation set for model selection and a fixed held-out test set. Our main results on the test set are presented in Table 1, and clinical evaluation is performed on this test set. In Table 1, we compare our ProtoVerse classification results with ProtoPNet and the post-hoc baseline model. For all the models, we used pre-trained VGG11 as the backbone and MWCE for classification loss since it was found to be a better fit for the imbalance in the dataset. Using the MWCE for classification, we could ensure

| Models                        | IBD | Class-avg<br>Accuracy | Class-avg<br>F1 | Sample<br>Accuracy |
|-------------------------------|-----|-----------------------|-----------------|--------------------|
| Baseline                      | ✗   | 72.46                 | 60.82           | 86.03              |
| ProtoPNet (Chen et al., 2019) | ✓   | 64.81                 | 51.59           | 83.24              |
| ProtoVerse (ours)             | ✓   | <b>76.28</b>          | <b>67.97</b>    | <b>90.22</b>       |

Table 1: Comparison of our ProtoVerse model with ProtoPNet and a baseline classifier on the test set. All models have been trained on Verse’19 dataset using MWCE for class imbalance. ProtoVerse, thanks to novel Diversity loss, significantly outperforms ProtoPNet and non-IBD baseline.

reasonably good class-average test accuracies even for such a highly imbalanced dataset. The class-average accuracy for each class for each of the models are given as - Baseline: [88.04, 72.22, 57.14]%, ProtoPNet: [86.5, 22.22, 85.71]%, ProtoVerse: [92.33, 72.22, 64.29]%. We observe that ProtoPNet does not perform as well as the baseline model in terms of the considered metrics. We attribute this to the feature collapse in the prototype layer. Importantly, we observed that the ProtoVerse model significantly outperforms both ProtoPNet and baseline across all metrics. For our ProtoVerse model, we found ( $\lambda_{\text{div}} = 0.3$ ) to be the optimum coefficient for Diversity loss. We postulate that explicitly ensuring diverse prototypes is beneficial to learn IBD models, especially in case of small sample size. The Diversity loss forces the network to focus on different parts of the vertebrae body, which is in line with clinical diagnosis rationale so that the variations in a fracture class can be adequately explained using the learnt prototypes.

| Models                        | IBD | Class-avg<br>Accuracy              | Class-avg<br>F1                    | Sample<br>Accuracy                 |
|-------------------------------|-----|------------------------------------|------------------------------------|------------------------------------|
| Baseline                      | ✗   | 70.92 $\pm$ 3.29                   | 63.52 $\pm$ 4.68                   | 89.26 $\pm$ 3.84                   |
| ProtoPNet (Chen et al., 2019) | ✓   | 72.73 $\pm$ 4.47                   | 67.27 $\pm$ 4.06                   | 91.13 $\pm$ 3.56                   |
| ProtoVerse (ours)             | ✓   | <b>76.36 <math>\pm</math> 5.96</b> | <b>75.97 <math>\pm</math> 4.80</b> | <b>93.58 <math>\pm</math> 1.37</b> |

Table 2: 5-fold cross-validation of our ProtoVerse model with ProtoPNet and a baseline classifier. Average and standard deviations over 5 runs are reported.

Additionally, we have performed a 5-fold cross-validation on the train set among non-IBD baseline, ProtoPNet, and ProtoVerse in Table 2. The general trend remains the same in the cross-validation as observed for the test set in Table 1. Accuracy for each class averaged across 5 splits are as follows: Baseline: [92.75, 39.99, 80]%, ProtoPNet: [93.89, 58.33, 66]%, ProtoVerse: [96.732, 62.33, 70]%. We also performed a pairwise Wilcoxon signed-rank test to determine the statistical significance between ProtoVerse, ProtoPNet, and Baseline. For Class-avg Accuracy, we find p-values 0.094 (ProtoVerse vs. Baseline) and 0.031 (ProtoVerse vs. ProtoPNet), respectively. For Class-avg F1, we find p-values 0.031 (ProtoVerse vs. Baseline) and 0.031 (ProtoVerse vs. ProtoPNet), respectively. For Sample Accuracy, we find p-values 0.031 (ProtoVerse vs. Baseline) and 0.031 (ProtoVerse vs. ProtoPNet),

respectively. ProtoVerse shows statistically significant ( $p\text{-value}<0.05$ ) results in two out of three metrics compared to baselines and all three metrics for ProtoPNet.




Figure 3: Cosine similarity between prototype vectors obtained from ProtoPNet and our ProtoVerse. Note that cosine similarity within a class is relatively high in the case of ProtoPNet, indicating difficulties in encompassing diverse prototypes. ProtoVerse achieves a diverse set of prototypes, as reflected in cosine similarity scores. Note that the third healthy prototype in both cases is looking into the trachea, which results in totally different features than vertebrae.

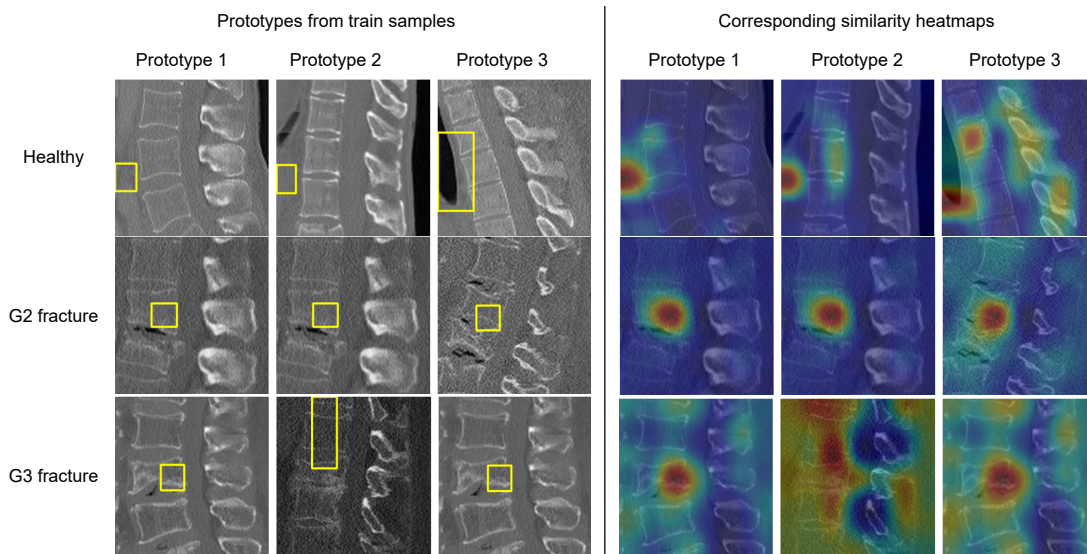
## 5.2 Prototype Quality of ProtoPNet vs ProtoVerse

We present prototypes from ProtoPNet and ProtoVerse models in Fig. 4 and cosine similarity among the prototype vectors in Fig. 3. We see that ProtoPNet produces repetitive prototypes for G2 and G3 and focuses outside the vertebral body for the healthy class. In contrast, ProtoVerse produces diverse prototypes for all three classes. Note that since most of the thoracic vertebrae do not have any fracture, both models see the nearby trachea region of those cases (3rd prototype of healthy class) as an attribute of healthy class.

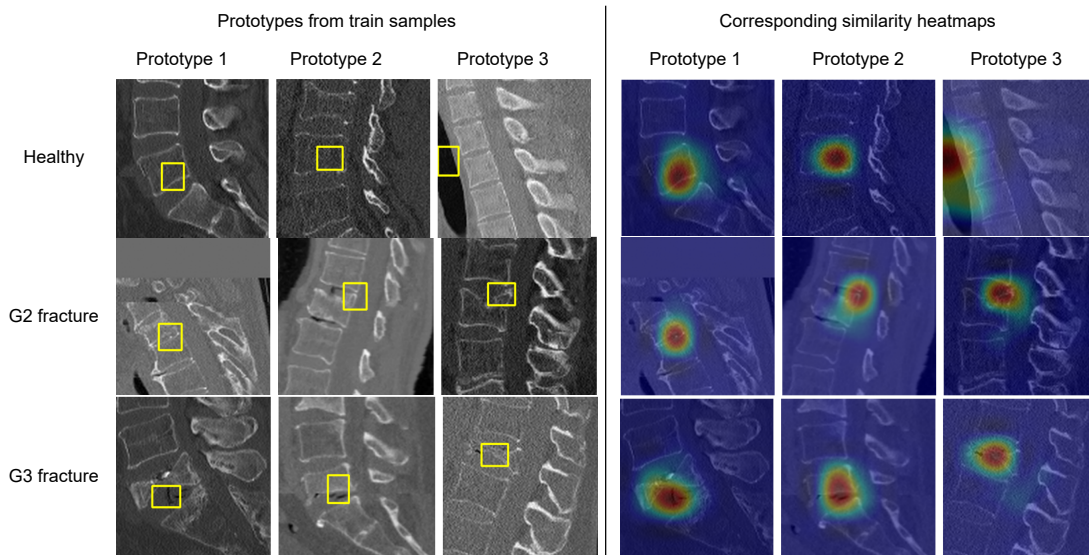
## 5.3 Visual Interpretability Analysis

We compare ProtoVerse’s interpretability with the post-hoc methods on the baseline model using Class Activation Map (CAM) in Fig. 5 and 6 for representative test samples from G2 and G3, respectively. Here, we compute the *similarity heatmap* (col. 2), *similarity score*, and *class connection score* as described in Section 4.1. The similarity score is multiplied by the class connection score to get the total contribution score of a given prototype to explain the classification output of a sample. We show the top 3 most activated prototypes corresponding to each of the test samples and the corresponding heatmap in col. 3 and col. 4, respectively. Col. 1 shows the highest activated region in the similarity heatmap for the test samples corresponding to each of the top-3 prototypes.

First sample in Fig. 5 has the first prototype focuses from top-right to mid-bottom regions of the vertebra, while the last two prototypes focuses on the mid-right sections. Second sample in Fig. 5 has the first two prototypes activating at the top-right corner of the vertebral body, while the last prototype attends to its middle section. The similarity



(a) ProtoPNet prototypes



(b) ProtoVerse prototypes

Figure 4: Qualitative comparison of prototype learnt from ProtoPNet and ProtoVerse model. Note that ProtoPNet produces repetitive prototypes whereas ProtoVerse captures diverse prototypes.

heatmaps focus on distinct relevant contexts around the deformities of G2. Positive class connection scores for all three prototypes from G2 confirm human-interpretable reasoning for classifying the test sample as G2.

Similarly, in both examples from Fig. 6, we observe that all the prototypes from G3 correspond to severely diminished height in the middle section of the vertebrae. These prototypes show high activations in similarly deformed sections at the middle and middle-

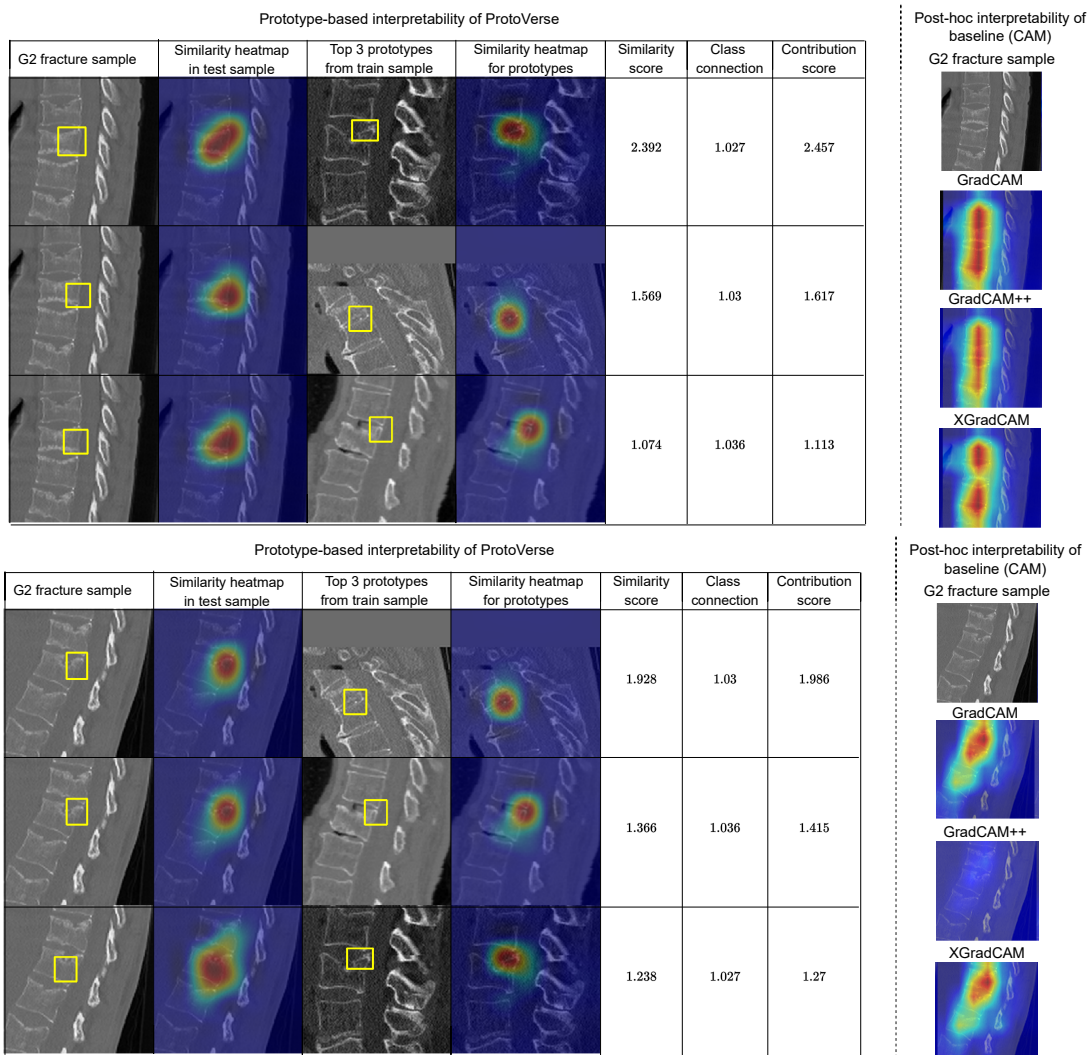


Figure 5: Test-time interpretability of ProtoVerse (left) in comparison to post-hoc baselines (right) for two typical G2 fracture samples. We show the top 3 closest prototypes (col. 3) based on similarity score and corresponding heatmaps (col. 4). Firstly, the similarity heatmap of ProtoVerse is localised to the fracture regions of the vertebrae (col 2). In contrast, post-hoc baselines fails to precisely localise clinically important regions in the vertebral body and focuses on the wrong vertebra. Secondly, the top 3 prototypes visually explain why the highest activated region is important to grade the fracture type correctly. Note that, for all 3 prototypes show positive class connections and belong to the same class.

right of the vertebrae. This observation reaffirms the reasoning ability of ProtoVerse based on meaningfully retrieved visual prototypes, e.g., similarly deformed vertebral edges. We also note that the last prototype (from G2) has a negative class connection since it does not belong to the class of the test sample despite having high visual similarity. Thus, the overall contribution score determines the role of the prototype in class prediction of the given sample. From both cases, one can observe the benefit of a diverse set of learnt

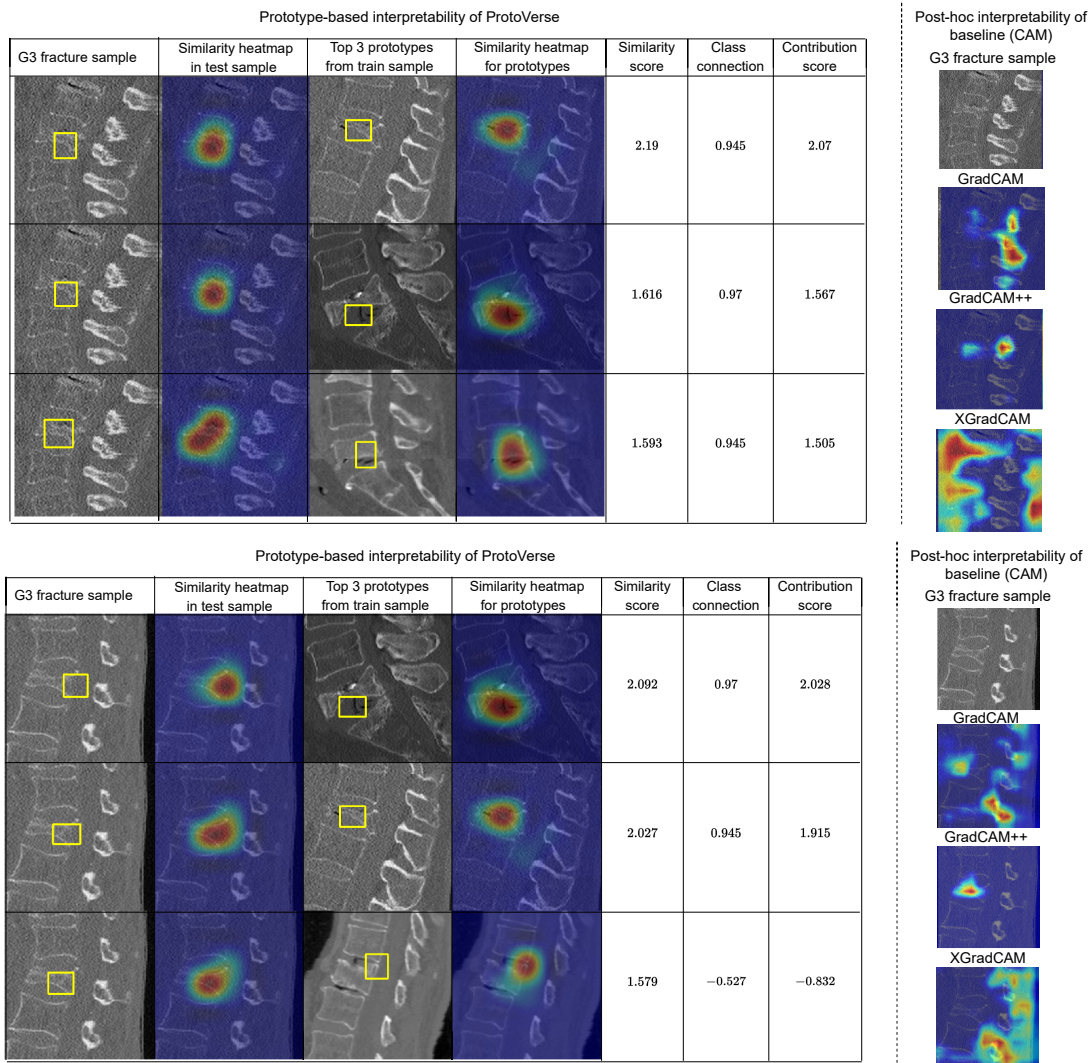


Figure 6: Test-time interpretability of ProtoVerse (left) in comparison to post-hoc baselines (right) for two typical G3 fracture samples. We show the top 3 closest prototypes (col. 3) based on similarity score and corresponding heatmaps (col. 4). Firstly, the similarity heatmap of ProtoVerse is localised to the fracture regions of the vertebrae (col 2). In contrast, post-hoc baselines fails to precisely localise clinically important regions in the vertebral body and focuses on the vertebral processes. Secondly, the top 3 prototypes visually explain why the highest activated region is important to grade the fracture type correctly. Note that, for first example, all 3 prototypes show positive class connections and belong to the same class, whereas the third prototype for second example coming from G2 shows a negative class connection.

prototypes in identifying different sub-parts of the fractured vertebrae of interest bringing in interpretability in the overall classification decision.

| Test Class | Prototype Class |       |       |
|------------|-----------------|-------|-------|
|            | Healthy         | G2    | G3    |
| Healthy    | 4.72            | -0.33 | -0.19 |
| G2         | 0.73            | 2.72  | -0.01 |
| G3         | 0.35            | -0.06 | 2.70  |

Table 3: The average contribution score for all samples from each class has been correlated with the prototypes from each class. Strong diagonal correlation indicates prototypes of each class are most activated and thus reliably explain samples from their class.

For post-hoc CAM methods, the heatmaps are generated a posteriori using the class activations from the baseline model. In our experiments, we considered three popular variants of CAM based methods, namely GradCAM ((Selvaraju et al., 2017)), GradCAM++ ((Chattopadhyay et al., 2017) and XGradCAM ((Fu et al., 2020)).

For the G2 samples, we observe that the most-activated (in red) is not highlighting the vertebrae of interest (center vertebrae) but focusing on the one above. For the G3 samples, the activation is scattered and does not focus on any vertebrae body for most of the methods except GradCAM++ for the second sample. Although the baseline model does not perform too badly in terms of classification accuracy, both cases indicate that it is barely interpretable post-hoc. This is because GradCAM variants only provides local interpretation based on sample-specific activation and ignores the recurring features important to explain the entire class, unlike the prototypes in ProtoVerse.

#### 5.4 Similarity-based Interpretability across Classes

Here, we analyse how much one class’s prototypes influence another class’s decision-making process. For that, we average total contribution scores from each class’s prototypes for all test samples of each class and present them in Table 3. We demonstrate that the average contribution for all the samples corresponding to the prototypes of their respective class is maximum (strong correlation along diagonal), thus confirming that the learnt prototypes in each class bear a strong correlation with the samples from the respective class. Additionally, the positive self-correlation of the healthy class (4.72) compared to negative cross-correlation in fracture classes G2, G3 ( $-0.33$ ,  $-0.19$ ) shows the model’s strong anomaly detection (healthy vs fracture) performance. To the best of our knowledge, our study is the first of its kind to successfully learn fine-grained prototypes to further distinguish between fracture grades beyond contemporary anomaly detection works (Mohammadjafari et al., 2021; Wei et al., 2023).

#### 5.5 Clinical Validation Results

Along with experimental results, we designed an expert clinical validation performed by two radiologists with 5 mean years of experience. For this, we emphasised the clinical usability of the learnt prototypes in explaining a test image. We prepared a study including all test samples of clinically relevant G2 and G3 fracture classes. For each sample, we listed the



top 3 closest prototypes from the training set based on the similarity score and asked the radiologists to rate based on the following two (Yes/No) questions (see Suppl. for details):

1. Are each of the retrieved prototypes relevant in explaining fracture grading of the test image, i.e., *Prototype Relevance*?
2. Are the highest activated parts of the test image visually similar to the highest activated parts of the retrieved prototypes, i.e., *Visual Similarity*?

| Expert Rater | Prototype Relevance |      |       | Visual Similarity |      |       |
|--------------|---------------------|------|-------|-------------------|------|-------|
|              | G2                  | G3   | Total | G2                | G3   | Total |
| Rater-1      | 1.00                | 1.00 | 1.00  | 0.33              | 0.72 | 0.59  |
| Rater-2      | 0.88                | 0.78 | 0.84  | 0.61              | 0.71 | 0.65  |

Table 4: Clinical Validation results show the high relevance of the learnt prototypes from both experts in explaining the test-time decision. Doctors found relatively higher visual similarity for G3 than G2, indicating difficulty in learning G2 from small datasets.

For both questions, we consider ‘Yes’ answer for a sample if 2 out of 3 prototypes are marked ‘Yes’ by the doctors. We report the percentage of ‘Yes’ outcome for all questions in Tab. 4. We observe that both radiologists find the retrieved prototypes to be relevant for adequately explaining the fracture class. G3 prototypes have a better average relevance score. We attribute this to the fact that G2 samples are not severely fractured and can often be semantically difficult to learn. Qualitatively, the doctors have found that the prototype’s activation patterns were more pronounced in the posterior elements of the vertebral body. This finding might be due to the algorithm’s proficiency in recognizing higher-grade fractures, which usually involve the posterior aspect of the vertebral body, compared to low-grade fractures of the anterior vertebral body. Radiologists rate consistently high visual similarity for G3, further validating the visual explanation. However, for G2, inter-rater variability is more pronounced, which is commonly observed for lower fracture grades (Buckens et al., 2013). Our future work will explore explicitly modeling inter-rater variability in the prototype layer to better learn a set of prototypes for small datasets.

## 5.6 Quantitative Ablation Study

In this section, we provide 5-fold cross-validation on the training set for hyper-parameter search and ablation on three aspects of our proposed ProtoVerse model. First, we study the effect of the number of prototypes per class. Thirdly, we compare our median-weighted class imbalance loss with different loss weighting strategies. Finally, we show ablation on the coefficient of the Diversity loss.

### 5.6.1 ABLATION ON NUMBER OF PROTOTYPES

Table 5 reports quantitative performance using varying numbers of prototypes per class for ProtoVerse with VGG11 backbone. We observe that three prototypes per class are optimal.

Since a single prototype per class severely hurts the interpretability of the model, we refrain from including it in the experiment.

| # Prototypes | Class-avg Accuracy | Class-avg F1      | Sample Accuracy  |
|--------------|--------------------|-------------------|------------------|
| 2            | $72.24 \pm 7.60$   | $66.22 \pm 9.80$  | $88.99 \pm 4.87$ |
| 3            | $76.36 \pm 5.95$   | $75.97 \pm 4.80$  | $93.58 \pm 1.37$ |
| 4            | $75.40 \pm 3.06$   | $70.07 \pm 5.82$  | $91.01 \pm 3.53$ |
| 5            | $74.67 \pm 7.33$   | $66.13 \pm 11.28$ | $87.34 \pm 7.35$ |

Table 5: Performance comparison by varying the number of prototypes per class

| $\lambda_{\text{div}}$ | Class-avg Accuracy | Class-avg F1     | Sample Accuracy  |
|------------------------|--------------------|------------------|------------------|
| 0.1                    | $75.66 \pm 3.89$   | $69.45 \pm 4.04$ | $91.10 \pm 1.10$ |
| 0.3                    | $76.36 \pm 5.96$   | $75.97 \pm 4.80$ | $93.58 \pm 1.37$ |
| 0.5                    | $75.12 \pm 2.10$   | $71.99 \pm 4.58$ | $92.84 \pm 2.06$ |

Table 6: Performance comparison for different weighting of Diversity loss.

### 5.6.2 ABLATION ON WEIGHTING STRATEGY

Fig. 7 shows ablation on different class weighting strategies to mitigate the class imbalance issue for VCF grading using three prototypes per class. Note that among all competing weights, median weighting produces the best result.




Figure 7: Bar plot showing comparison among different class weighting strategies

### 5.6.3 ABLATION ON DIVERSITY LOSS COEFFICIENT

Table 6 shows the ablation of coefficient for Diversity loss. We use the above optimal parameters for the model. We note that 0.3 is a suitable choice for sufficient variation in the visual prototype without affecting the model’s classification performance.

### Limitation

We believe improving the dataset imbalance with an emphasis on diverse fracture grade samples (potentially G1 fractures) could further improve the quality of the prototypes and clinical utility of such interpretable models. Lastly, although the learnt prototypes are designed to be human-understandable, nonetheless some prototypes are not human-interpretable and it still remains to be explored whether the model reasoning is incorrect in such cases or the generated explanation needs to be enhanced to be better human-aligned.

## 6. Conclusion

In this work, we proposed a novel IBD method, ProtoVerse, for VCF grading, leveraging prototype-based representation learning. Specifically, we address the challenge of learning diverse yet representative prototypes from a small dataset with a high class-imbalance. The resultant ProtoVerse outperforms previous IBD models, such as ProtoPNet, and offers enhanced interpretability that non-IBD baselines. We have validated our results with radiologists, confirming the clinical applicability of the proposed method. To further enhance the quality of learnt prototypes, future research could include human-in-loop feedback during training to identify and replace unsatisfactory prototypes.

## Acknowledgments

This work was funded by the Bavarian Ministry for Economic Affairs, Regional Development and Energy as part of a project to support the thematic development of the Institute for Cognitive Systems (IKS). SS and JSK are supported by European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (101045128-iBack-epic-ERC2021-COG).

## Ethical Standards

The work follows appropriate ethical standards in conducting research and writing the manuscript, following all applicable laws and regulations regarding treatment of animals or human subjects.

## Conflicts of Interest

We declare we don’t have conflicts of interest.

## References

- Nishanth Arun, Nathan Gaw, Praveer Singh, Ken Chang, Mehak Aggarwal, Bryan Chen, Katharina Hoebel, Sharut Gupta, Jay Patel, Mishka Gidwani, et al. Assessing the trustworthiness of saliency maps for localizing abnormalities in medical imaging. *Radiology: Artificial Intelligence*, 3(6):e200267, 2021.
- Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On Pixel-Wise Explanations for Non-Linear Classifier Decisions by Layer-Wise Relevance Propagation. *PLOS ONE*, 10(7):e0130140, July 2015. ISSN 1932-6203. .
- Ghada Ballane, JA Cauley, MM Luckey, and G El-Hajj Fuleihan. Worldwide prevalence and incidence of osteoporotic vertebral fractures. *Osteoporosis International*, 28:1531–1542, 2017.
- Amir Bar, Lior Wolf, Orna Bergman Amitai, Eyal Toledano, and Eldad Elnekave. Compression fractures detection on ct. In *Medical imaging 2017: computer-aided diagnosis*, volume 10134, pages 1036–1043. SPIE, 2017.
- Constantinus F Buckens, Pim A de Jong, Christian Mol, Eric Bakker, Hein P Stallman, Willem P Mali, Yolanda van der Graaf, and Helena M Verkooijen. Intra and interobserver reliability and agreement of semiquantitative vertebral fracture assessment on chest computed tomography. *PloS one*, 8(8):e71204, 2013.
- Kaidi Cao, Maria Brbic, and Jure Leskovec. Concept Learners for Few-Shot Learning. *arXiv:2007.07375 [cs, stat]*, March 2021.
- George A Carberry, B Dustin Pooler, Neil Binkley, Travis B Lauder, Richard J Bruce, and Perry J Pickhardt. Unreported vertebral body compression fractures at abdominal multidetector ct. *Radiology*, 268(1):120–126, 2013.
- Gianluca Carloni, Andrea Berti, Chiara Iacconi, Maria Antonietta Pascali, and Sara Colantonio. On the applicability of prototypical part learning in medical images: breast masses classification using protopnet. In *International Conference on Pattern Recognition*, pages 539–557. Springer, 2022.
- Zachariah Carmichael, Suhas Lohit, Anoop Cherian, Michael J Jones, and Walter J Scheirer. Pixel-grounded prototypical part networks. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 4768–4779, 2024.
- JA Cauley, DE Thompson, KC Ensrud, JC Scott, and D Black. Risk of mortality following clinical fractures. *Osteoporosis international*, 11:556–561, 2000.
- Aditya Chattopadhyay, Anirban Sarkar, Prantik Howlader, and Vineeth N. Balasubramanian. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 839–847, 2017. URL <https://api.semanticscholar.org/CorpusID:13678776>.

- Chaofan Chen, Oscar Li, Daniel Tao, Alina Barnett, Cynthia Rudin, and Jonathan K Su. This Looks Like That: Deep Learning for Interpretable Image Recognition. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- David Chettrit, Tomer Meir, Hila Lebel, Mila Orlovsky, Ronen Gordon, Ayelet Akselrod-Ballin, and Amir Bar. 3d convolutional sequence to sequence model for vertebral compression fractures identification in ct. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part VI 23*, pages 743–752. Springer, 2020.
- Jon Donnelly, Alina Jade Barnett, and Chaofan Chen. Deformable ProtoPNet: An Interpretable Image Classifier Using Deformable Prototypes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10265–10275, 2022.
- Paul Engstler, Matthias Keicher, David Schinz, Kristina Mach, Alexandra S. Gersing, Sarah C. Foreman, Sophia S. Goller, Juergen Weissinger, Jon Rischewski, Anna-Sophia Dietrich, Benedikt Wiestler, Jan S. Kirschke, Ashkan Khakzar, and Nassir Navab. Interpretable Vertebral Fracture Diagnosis. In Mauricio Reyes, Pedro Henriques Abreu, and Jaime Cardoso, editors, *Interpretability of Machine Intelligence in Medical Image Computing*, Lecture Notes in Computer Science, pages 71–81, Cham, 2022a. Springer Nature Switzerland. ISBN 978-3-031-17976-1. .
- Paul Engstler, Matthias Keicher, David Schinz, Kristina Mach, Alexandra S Gersing, Sarah C Foreman, Sophia S Goller, Juergen Weissinger, Jon Rischewski, Anna-Sophia Dietrich, et al. Interpretable vertebral fracture diagnosis. In *Interpretability of Machine Intelligence in Medical Image Computing: 5th International Workshop, iMIMIC 2022, Held in Conjunction with MICCAI 2022, Singapore, Singapore, September 22, 2022, Proceedings*, pages 71–81. Springer, 2022b.
- Zhengqing Fang, Kun Kuang, Yuxiao Lin, Fei Wu, and Yu-Feng Yao. Concept-based Explanation for Fine-grained Images and Its Application in Infectious Keratitis Classification. In *Proceedings of the 28th ACM International Conference on Multimedia*, MM '20, pages 700–708, New York, NY, USA, October 2020. Association for Computing Machinery. ISBN 978-1-4503-7988-5. .
- Shixiang Feng, Beibei Liu, Ya Zhang, Xiaoyun Zhang, and Yuehua Li. Two-stream compare and contrast network for vertebral compression fracture diagnosis. *IEEE Transactions on Medical Imaging*, 40(9):2496–2506, 2021.
- Ruth Fong and Andrea Vedaldi. Net2Vec: Quantifying and Explaining How Concepts are Encoded by Filters in Deep Neural Networks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8730–8738, June 2018. .
- Ruigang Fu, Qingyong Hu, Xiaohu Dong, Yulan Guo, Yinghui Gao, and Biao Li. Axiom-based grad-cam: Towards accurate visualization and explanation of cnns. *ArXiv*, abs/2008.02312, 2020. URL <https://api.semanticscholar.org/CorpusID:221006223>.

- Srishti Gautam, Marina M.-C. Höhne, Stine Hansen, Robert Jenssen, and Michael Kampffmeyer. This looks more like that: Enhancing Self-Explaining Models by Prototypical Relevance Propagation. *arXiv:2108.12204 [cs]*, August 2021.
- Harry K Genant, Chun Y Wu, Cornelis Van Kuijk, and Michael C Nevitt. Vertebral fracture assessment using a semiquantitative technique. *Journal of bone and mineral research*, 8 (9):1137–1148, 1993.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Adrian Hoffmann, Claudio Fanconi, Rahul Rade, and Jonas Kohler. This looks like that... does it? shortcomings of latent space prototype interpretability in deep networks. *arXiv preprint arXiv:2105.02968*, 2021.
- Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- Qihan Huang, Mengqi Xue, Wenqi Huang, Haofei Zhang, Jie Song, Yongcheng Jing, and Mingli Song. Evaluation and improvement of interpretability for self-explainable part-prototype networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2011–2020, 2023.
- Malek Husseini, Anjany Sekuboyina, Maximilian Loeffler, Fernando Navarro, Bjoern H. Menze, and Jan S. Kirschke. Grading Loss: A Fracture Grade-Based Metric Loss for Vertebral Fracture Detection. In Anne L. Martel, Purang Abolmaesumi, Danail Stoyanov, Diana Mateus, Maria A. Zuluaga, S. Kevin Zhou, Daniel Racoceanu, and Leo Joskowicz, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, Lecture Notes in Computer Science, pages 733–742, Cham, 2020. Springer International Publishing. .
- Dmitry Kazhdan, Boty Dimanov, Mateja Jamnik, Pietro Li, and Adrian Weller. Now You See Me (CME): Concept-based Model Extraction. page 15.
- Matthias Keicher, Matan Atad, David Schinz, Alexandra S. Gersing, Sarah C. Foreman, Sophia S. Goller, Juergen Weissinger, Jon Rischewski, Anna-Sophia Dietrich, Benedikt Wiestler, Jan S. Kirschke, and Nassir Navab. Semantic Latent Space Regression of Diffusion Autoencoders for Vertebral Fracture Grading, March 2023.
- Eunji Kim, Siwon Kim, Minji Seo, and Sungroh Yoon. XProtoNet: Diagnosis in Chest Radiography with Global and Local Explanations. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15714–15723, Nashville, TN, USA, June 2021. IEEE. ISBN 978-1-66544-509-2. .
- Pieter-Jan Kindermans, Sara Hooker, Julius Adebayo, Maximilian Alber, Kristof T Schütt, Sven Dähne, Dumitru Erhan, and Been Kim. The (un) reliability of saliency methods.

- Explainable AI: Interpreting, explaining and visualizing deep learning*, pages 267–280, 2019.
- Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. Concept Bottleneck Models. *arXiv:2007.04612 [cs, stat]*, December 2020.
- Oscar Li, Hao Liu, Chaofan Chen, and Cynthia Rudin. Deep Learning for Case-Based Reasoning Through Prototypes: A Neural Network That Explains Its Predictions. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), April 2018. ISSN 2374-3468.
- Maximilian T Löffler, Anjany Sekuboyina, Alina Jacob, Anna-Lena Grau, Andreas Scharr, Malek El Husseini, Mareike Kallweit, Claus Zimmer, Thomas Baum, and Jan S Kirschke. A vertebral segmentation dataset with fracture grading. *Radiology: Artificial Intelligence*, 2(4):e190138, 2020.
- Sanaz Mohammadjafari, Mucahit Cevik, Mathusan Thanabalasingam, Ayse Basar, and Alzheimer’s Disease Neuroimaging Initiative. Using ProtoPNet for Interpretable Alzheimer’s Disease Classification. *Proceedings of the Canadian Conference on Artificial Intelligence*, June 2021. .
- Meike Nauta, Ron van Bree, and Christin Seifert. Neural Prototype Trees for Interpretable Fine-grained Image Recognition. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14928–14938, Nashville, TN, USA, June 2021. IEEE. ISBN 978-1-66544-509-2. .
- Anh Nguyen, Jason Yosinski, and Jeff Clune. Understanding Neural Networks via Feature Visualization: A Survey. In Wojciech Samek, Grégoire Montavon, Andrea Vedaldi, Lars Kai Hansen, and Klaus-Robert Müller, editors, *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, Lecture Notes in Computer Science, pages 55–76. Springer International Publishing, Cham, 2019. ISBN 978-3-030-28954-6. .
- Joeri Nicolaes, Steven Raeymaeckers, David Robben, Guido Wilms, Dirk Vandermeulen, Cesar Libanati, and Marc Debois. Detection of vertebral fractures in ct using 3d convolutional neural networks. In *Computational Methods and Clinical Applications for Spine Imaging: 6th International Workshop and Challenge, CSI 2019, Shenzhen, China, October 17, 2019, Proceedings 6*, pages 3–14. Springer, 2020.
- Maxim Pisov, Vladimir Kondratenko, Alexey Zakharov, Alexey Petraikin, Victor Gombolevskiy, Sergey Morozov, and Mikhail Belyaev. Keypoints localization for joint vertebra detection and fracture severity quantification. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part VI 23*, pages 723–732. Springer, 2020.
- Amy Rafferty, Rudolf Nenutil, and Ajitha Rajan. Explainable artificial intelligence for breast tumour classification: Helpful or harmful. In *International Workshop on Interpretability of Machine Intelligence in Medical Image Computing*, pages 104–123. Springer, 2022.

- Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.
- Dawid Rymarczyk, Łukasz Struski, Michał Górszczak, Koryna Lewandowska, Jacek Tabor, and Bartosz Zieliński. Interpretable Image Classification with Differentiable Prototypes Assignment, September 2022.
- Wojciech Samek and Klaus-Robert Müller. Towards Explainable Artificial Intelligence. In Wojciech Samek, Grégoire Montavon, Andrea Vedaldi, Lars Kai Hansen, and Klaus-Robert Müller, editors, *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, Lecture Notes in Computer Science, pages 5–22. Springer International Publishing, Cham, 2019. .
- Adriel Saporta, Xiaotong Gui, Ashwin Agrawal, Anuj Pareek, Steven QH Truong, Chanh DT Nguyen, Van-Doan Ngo, Jayne Seekins, Francis G Blankenberg, Andrew Y Ng, et al. Benchmarking saliency methods for chest x-ray interpretation. *Nature Machine Intelligence*, 4(10):867–878, 2022.
- Gesina Schwalbe. Concept Embedding Analysis: A Review. *arXiv:2203.13909 [cs, stat]*, March 2022.
- Anjany Sekuboyina, Malek E Hussein, Amirhossein Bayat, Maximilian Löffler, Hans Liebl, Hongwei Li, Giles Tetteh, Jan Kukačka, Christian Payer, Darko Štern, et al. Verse: a vertebrae labelling and segmentation benchmark for multi-detector ct images. *Medical image analysis*, 73:102166, 2021.
- Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Gurmail Singh and Kin-Choong Yow. These do not look like those: An interpretable deep learning model for image recognition. *IEEE Access*, 9:41482–41493, 2021.
- Poulami Sinhamahapatra., Lena Heidemann., Maureen Monnet., and Karsten Roscher. Towards human-interpretable prototypes for visual assessment of image classification models. In *Proceedings of the 18th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2023) - Volume 5: VISAPP*, pages 878–887. INSTICC, SciTePress, 2023. ISBN 978-989-758-634-7. .
- Susu Sun, Lisa M Koch, and Christian F Baumgartner. Right for the wrong reason: Can interpretable ml techniques detect spurious correlations? In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 425–434. Springer, 2023.



- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic Attribution for Deep Networks. In *Proceedings of the 34th International Conference on Machine Learning*, pages 3319–3328. PMLR, July 2017.
- Erico Tjoa and Cuntai Guan. A Survey on Explainable Artificial Intelligence (XAI): Towards Medical XAI. *arXiv:1907.07374 [cs]*, August 2020. .
- Naofumi Tomita, Yvonne Y Cheung, and Saeed Hassanpour. Deep neural networks for automatic detection of osteoporotic vertebral fractures on ct scans. *Computers in biology and medicine*, 98:8–15, 2018.
- Chong Wang, Yuyuan Liu, Yuanhong Chen, Fengbei Liu, Yu Tian, Davis McCarthy, Helen Frazer, and Gustavo Carneiro. Learning support and trivial prototypes for interpretable image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2062–2072, 2023.
- Yuanyuan Wei, Roger Tam, and Xiaoying Tang. MProtoNet: A Case-Based Interpretable Model for Brain Tumor Classification with 3D Multi-parametric Magnetic Resonance Imaging, April 2023.
- Eren Bora Yilmaz, Christian Buerger, Tobias Fricke, Md Motiur Rahman Sagar, Jaime Peña, Cristian Lorenz, Claus-Christian Glüer, and Carsten Meyer. Automated deep learning-based detection of osteoporotic fractures in ct images. In *Machine Learning in Medical Imaging: 12th International Workshop, MLMI 2021, Held in Conjunction with MICCAI 2021, Strasbourg, France, September 27, 2021, Proceedings 12*, pages 376–385. Springer, 2021.
- Alexey Zakharov, Maxim Pisov, Alim Bukharaev, Alexey Petraikin, Sergey Morozov, Victor Gombolevskiy, and Mikhail Belyaev. Interpretable vertebral fracture quantification via anchor-free landmarks localization. *Medical Image Analysis*, 83:102646, 2023.
- Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning Deep Features for Discriminative Localization. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2921–2929, June 2016. .

### Appendix A. Details on Clinical Validation

Fig. 8 shows a sample from the clinical validation report. The radiologists are provided with the top-3 activated prototypes along with the highest-activated regions in the test image. Additionally, they are given the similarity heatmap of each of the top-3 prototypes for visual reasoning. Both radiologists provided ‘yes/no ’answers to the questions and overall opinion on the findings.

A sample from Clinical Validation study


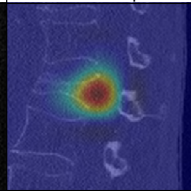

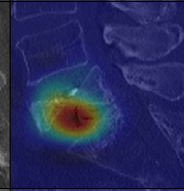

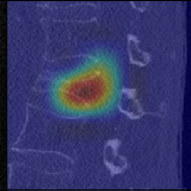

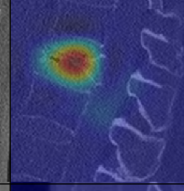

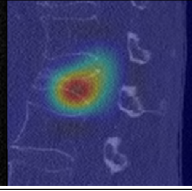

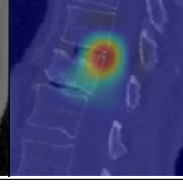
| G3 fracture sample   | Similarity heatmap in test sample  | Top 3 prototypes from train sample   | Similarity heatmap for prototypes   | Prototype class id | Prototype Relevance | Visual Similarity |
|--|--|--|---|--------------------|---------------------|-------------------|
|   |   |   |   | 2                  | Yes,<br>Yes         | No,<br>Yes        |
|   |   |   |   | 2                  | Yes,<br>Yes         | Yes,<br>Yes       |
|  |  |  |  | 1                  | Yes,<br>No          | No,<br>No         |

Figure 8: Exemplary samples from clinical evaluation done by radiologists. The answers obtained from the two radiologists are indicated.