

ProtoMedX: Towards Explainable Multi-Modal Prototype Learning for Bone Health Classification

Alvaro Lopez Pellicer¹ André Mariucci¹ Plamen Angelov¹
Marwan Bukhari² Jemma G. Kerns²

¹School of Computing and Communications; ²Lancaster Medical School

Lancaster University, Lancaster, United Kingdom

{a.lopezpellicer, a.mariucci, p.angelov, m.bukhari, j.kerns}@lancaster.ac.uk

Abstract

*Bone health studies are crucial in medical practice for the early detection and treatment of Osteopenia and Osteoporosis. Clinicians usually make a diagnosis based on densitometry (DEXA scans) and other patient history. The applications of AI in this field are an ongoing research. Most of the successful methods for this task include Deep Learning models that rely on vision alone (DEXA / X-ray imagery) geared towards high prediction accuracy, where explainability is disregarded and largely based on the post hoc assessment of input contributions. We propose **ProtoMedX**, a multi-modal model that uses both DEXA scans of the lumbar spine and patient records. ProtoMedX's prototype-based architecture is explainable by design, crucial for medical applications, especially in the context of the upcoming EU AI Act, as it allows explicit analysis of the model's decisions, especially the ones that are incorrect. ProtoMedX demonstrates state-of-the-art performance in bone health classification while also providing explanations that can be visually understood by clinicians. Using our dataset of 4,160 real NHS patients, the proposed ProtoMedX achieves 87.58% accuracy in vision-only tasks and 89.8% in its multi-modal variant, both approaches surpassing existing published methods.*

1. Introduction

In 2022, osteoporosis affected 3.5 million people in the United Kingdom, with osteoporotic fractures accounting for over £4.6 billion of direct costs to the National Health Service (NHS). This annual cost is forecasted to rise to £6 billion by 2030 as the population of the United Kingdom ages [24]. Globally, the prevalence of osteoporosis is estimated to be 6.3% of men over the age of 50 and 21.2% of women over the same age range [10]. Based on the world population, this suggests that approximately 500 million individu-

als worldwide may be affected [10].

Existing bone health classification approaches face three fundamental limitations. First, most studies adopt binary classification (normal vs. osteoporosis), either excluding osteopenia patients altogether or grouping them with osteoporosis cases under the label “low bone density” [32]. This simplification obscures the clinically important intermediate state of osteopenia, where patients face a significantly higher fracture risk compared to those with osteoporosis [30]. Possibly due to data availability issues, studies attempting three-class classification (normal, osteopenia, and osteoporosis) are much rarer. The only three-class study we identified using spine data, [36], did not report classification accuracy but achieved a highest AUC of 0.81, suggesting that the added complexity poses substantial challenges.

Second, current methods predominantly rely on vision-only models [14, 28]. This contradicts standard medical practice for fracture risk assessment, where diagnosis integrates image findings with patient history, demographics, and risk factors. Our results validate this intuition; by incorporating clinical features alongside imaging, accuracy improves from **87.58% to 89.8%** while providing more robust predictions.

Third, and critically, existing deep learning approaches lack explainability. *Post hoc* interpretation methods such as GradCAM [27] and SHAP values [15] provide approximate rationalisations post-training rather than true model reasoning. In medical contexts where decisions directly impact patient care, clinicians require models that are explainable by design and cannot rely on “black boxes” with retrofitted interpretations to guide their treatment decisions [25]. This importance is further emphasised in the context of the recently adopted EU AI Act, which considers medical applications of AI as “high risk”, making it as significant as critical infrastructure [7, 7, 19].

To address these limitations, we present **ProtoMedX** (**Proto**type-based **Med**ical **eX**planation), the first

prototype-learning framework for bone health assessment. ProtoMedX reimagines bone health classification through case-based reasoning, rather than learning opaque decision boundaries. It identifies representative prototypes for each diagnostic category and classifies new patients based on similarities to these learned examples. This mirrors actual clinical reasoning, where physicians compare patients to archetypal cases. Our key contributions are:

- **First prototype-based architecture for bone health:** We introduce dual prototype spaces (visual and clinical) unified through cross-modal attention, enabling explainable predictions via direct comparison to learned exemplars. (See Figures 2, 5).
- **Multi-task learning leveraging bone density continuity:** We demonstrate that auxiliary T-score regression during training forces the model to understand bone density as a continuous phenomenon, dramatically improving classification accuracy by 6.7% in MLP approaches (Table 1) and by 2.46% in ProtoMedX (See Table 2).
- **State-of-the-art performance with built-in explainability:** ProtoMedX achieves 89.8% accuracy on three-class classification, surpassing prior art by 14-27% absolute improvement (Table 4). Critically, this performance includes inherent explanations, not **post hoc** rationalisations, making it more suitable for clinical deployment.
- **Comprehensive evaluation on clinical data:** The study is comprised of 4,160 lumbar spine DEXA scans with complete patient profiles, demonstrating robust performance across demographics. The model excels at clinically critical Normal vs. Abnormal detection (91.2% sensitivity) while maintaining prototype-based transparency. (See Sec. 4.1).

2. Background and Related Work

2.1. Bone Health

As we age, we naturally lose bone mass. This happens to some patients at a faster rate than others, which can result in osteoporosis. Osteoporosis is a bone disease that results in an increased risk of a fragility fracture, which occurs following a fall from standing height or less. Osteoporosis affects both men and women but is most prevalent in postmenopausal women due to reduced oestrogen, with 1 in 2 women affected by the age of 80 [10]. While it is not currently possible to completely recover bone after significant degradation, it is possible to reduce further bone loss and the impact through lifestyle and/or pharmaceutical interventions if detected sufficiently early. There are several clinical tools available to assess fracture risk, most notably FRAX® [11], which calculates a patient’s probability of suffering from a hip fracture caused by osteoporosis within 10 years, as well as provides advice as to whether a patient should seek a bone densitometry scan produced by a DEXA ma-

chine [17]. The DEXA machine (Dual-Energy X-ray Absorptiometry) uses a low dose of X-rays to measure bone mineral density and provides a T-score. The T-score indicates the number of standard deviations away a patient is from a young, healthy population and provides the diagnostic definition of bone health, according to the World Health Organization (WHO), using the following thresholds [23]: normal (T-score above -1.0), osteopenia (T-score between -1.0 and -2.5), or osteoporosis (T-score below -2.5) [23].

DEXA scans are commonly performed on hips, forearms, lumbar spine, and whole-body scans. For this study, we were granted the use of lumbar spine scans only.

2.2. Dataset Availability

Moderate-sized tabular datasets (1,000–10,000 samples) are available, such as [14] who employed a sample of 6,672 patients from an NHANES-provided dataset [3]. For vision tasks, publicly available datasets for use in research are often significantly smaller, normally in the hundreds, which limits the amount of diversity for each dataset and limits the quality of research by third-party institutions. One example of a publicly available dataset includes 177 lumbar DEXA scans from patients in Pakistan [16]. These public datasets are usually limited in size due to strict data restrictions, often resulting in datasets remaining proprietary. To the best of our knowledge, there are no large or moderate sized datasets that contain both images and more than a few patient attributes (outside of age, weight, height, and gender) that are available to the public, making multi-modal machine learning largely inaccessible. While our dataset is not currently accessible to the public, we support ongoing efforts to promote the sharing and accessibility of such datasets within the research community, provided the necessary data protection procedures are correctly implemented. Due to the aforementioned limitations in data availability, many studies [13, 32, 36] rely on data augmentation methods to resolve imbalanced data, which is standard for vision tasks. However, there are doubts on the transferability to real clinical practice application of models that leverage augmentation within the medical domain [20, 21, 31].

Our research was performed in collaboration with the NHS, which allowed us to obtain a significantly larger dataset than is common in this field. The dataset contains 4,160 patients’ lumbar spine DEXA scans, of which 2,224 are patients with normal bone health, 1,398 are patients with osteopenia, and 538 are patients with osteoporosis. The data consists of 926 males and 3,234 females. The ages of patients in the dataset range from 20 to 110, with an average age of 74 and an interquartile range of 65 to 83. The dataset also contains parts of the patients records, specifically, the same 11 features that the FRAX® tool operates with.

2.3. Machine Learning for Bone Health

Machine learning approaches to bone health classification have evolved from early neural networks [18, 22] to modern deep learning architectures. However, fundamental limitations persist across these methods.

Tabular-only approaches, such as those using demographic and bone turnover markers (e.g., age, sex, and bone-specific alkaline phosphatase), can achieve up to 73% accuracy [2]. However, these methods lack the spatial information provided by imaging.

Vision-only methods using CNNs and Vision Transformers on DEXA/X-ray images report 80-90% accuracy [14, 26, 28, 32]. However, these studies suffer from (1) binary classification that ignores the clinically important osteopenia stage [32], (2) small datasets (e.g., 117 images [32]) with heavy augmentation, (3) uncommon body areas, such as [26] using knees, which is not commonly used for osteoporosis diagnosis, or (4) poor three-class performance [36].

Multi-modal fusion shows modest improvements. [28] achieved 84.5% (vs. 83.2% vision-only) by combining dental X-rays with demographics. While this study employed binary classification, it demonstrates that the addition of tabular features in a multi-modal setting can provide increased accuracy, which is in line with our own findings.

These approaches share critical gaps: oversimplified binary classification, lack of explainability, and limited multi-modal integration, motivating our prototype-based solution that addresses all three.

2.4. Explainable AI for Bone Health

Clinical deployment demands models that are able to explain their reasoning, not just achieve high accuracy. *Post hoc* methods such as GradCAM [27], SHAP [15], and saliency maps [35] approximate model behaviour after training but lack fidelity and have the capability of being misleading, as they provide rationalisations rather than true model reasoning [8, 9, 12, 25, 34]. This disconnect is particularly problematic in medical contexts where understanding the decision process guides treatment.

Prototype-based Models: Prototype learning offers inherent explainability by comparing inputs to learned exemplars. ProtoPNet [4] pioneered this approach for images, followed shortly by xDNN [1], enabling case-based reasoning (“this looks like that”). However, ProtoPNet suffers from three limitations: (1) vision-only architecture makes the framework incompatible with clinical data, (2) reliance on localised patches that can highlight misleading regions,

and (3) *post hoc* similarity heatmaps rather than supervised prototype learning [29] (see Figure 1).




Figure 1. Inconsistent quality of prototypes and heatmap localisation generated by ProtoPNet. Despite (1a) showing a clear, high-quality prototype, its heatmap (1b) exhibits poor localisation with diffuse activation. Conversely, a blurry prototype (2a) produces well-focused heatmap localisation (2b), revealing that prototype visual quality does not correlate with localisation accuracy.

To the best of our knowledge, no prior work applies prototype-based learning to bone health classification. **ProtoMedX** addresses this gap by extending prototype reasoning to multi-modal inputs (imaging + clinical data) while learning full-image prototypes that provide clear explanations. This design aligns with the way clinicians reason, comparing patients to archetypal cases rather than analysing isolated image regions.

3. ProtoMedX

ProtoMedX introduces prototype-based deep learning for bone health assessment, enabling case-based reasoning and explanations that align with clinical decision-making. It learns class-representative prototypes in a unified feature space, making predictions by comparing inputs to these prototypes. Fig. 2 illustrates the overall architecture of the proposed framework.

3.1. Feature Extraction and Encoding

Vision Branch: We employ a frozen CrossViT backbone [5] pre-trained on ImageNet for DEXA feature extraction, as it achieved superior standalone performance (See table 3). CrossViT’s dual-branch design captures both fine-grained trabecular patterns and global spinal morphology, critical for osteoporosis assessment. The frozen backbone outputs $\mathbf{f}^{\text{img}} \in \mathbb{R}^{1151}$, which is fed to a projection MLP




Figure 2. **Overview of ProtoMedX Architecture.** Multi-modal prototype learning combines patient DEXA scans and clinical records *via* separate encoders, learns explainable vision and tabular prototypes, and fuses them in a joint prototype space. Classification and explanations derive from prototype similarity and case retrieval, enhancing clinical explainability.

(1151 $\rightarrow$ 512 $\rightarrow$ 256) with BatchNorm and 0.3 dropout, producing

$$\mathbf{h}_i = \phi_i(\mathbf{f}^{\text{img}}) \in \mathbb{R}^{256}.$$

Tabular Branch: We encode the clinical features as a vector $\mathbf{x}^{\text{tab}} \in \mathbb{R}^{11}$ (Age, sex, weight, height, previous fracture, parent fracture hip, currently smoking, glucocorticoids, rheumatoid arthritis, secondary osteoporosis, and alcohol intake of 3 or more units/day) *via*

$$\mathbf{h}_t = \phi_t(\mathbf{x}^{\text{tab}}), \quad \phi_t: \mathbb{R}^{11} \rightarrow \mathbb{R}^{64},$$

where ϕ_t is a two-layer MLP with residual connections and dropout.

3.2. Cross-Modal Fusion

The encoded image and tabular embeddings are first combined by a cross-modal attention block, producing a joint representation:

$$\mathbf{h}_{\text{fused}} = \text{CrossAtt}(\mathbf{h}_i, \mathbf{h}_t) \in \mathbb{R}^{256}.$$

During prototype matching, we apply an adaptive, input-dependent weighting of modality-specific similarity scores:

$$\alpha = \sigma(g([\mathbf{h}_i; \mathbf{h}_t])), \quad \text{sim}(\mathbf{z}) = \alpha S_{\text{img}}(\mathbf{z}) + (1 - \alpha) S_{\text{tab}}(\mathbf{z}),$$

where g is a linear layer and S_{img} , S_{tab} are cosine similarities in the image and tabular prototype subspaces, respectively. This gating strategy improved top-1 accuracy by 1.08% compared with uniform weighting. The fused vector $\mathbf{h}_{\text{fused}}$ is subsequently fed to the prototype layer for classification.

3.3. Prototype Learning

ProtoMedX learns $K = 6$ prototypes per class, where each prototype represents a typical patient case. This value was determined through empirical validation: $K < 6$ failed to

capture intra-class diversity (e.g., early vs. advanced osteoporosis), while $K > 6$ introduced redundant prototypes without improving classification. We employ a dual prototype architecture, maintaining separate spaces for visual ($\mathbf{p}_{c,k}^{\text{img}} \in \mathbb{R}^{128}$) and tabular ($\mathbf{p}_{c,k}^{\text{tab}} \in \mathbb{R}^{64}$) features while performing classification in the fused space. This enables modality-specific interpretation. Ablation shows that removing dual prototypes reduces accuracy by 4.33%.

Following contrastive learning principles [6], we optimise prototypes through a composite loss:

$$\mathcal{L}_{\text{proto}} = \mathcal{L}_{\text{class}} + \lambda_{\text{sep}} \mathcal{L}_{\text{sep}} + \lambda_{\text{ctr}} \mathcal{L}_{\text{center}} \quad (1)$$

where $\mathcal{L}_{\text{class}}$ ensures prototypes alone can classify *via* softmax over cosine similarities, $\mathcal{L}_{\text{sep}}$ maintains inter-class margins using a triplet-like loss [37], and $\mathcal{L}_{\text{center}}$ compacts intra-class distributions [33]. This creates well-separated clusters (Fig. 3) with osteoporosis prototypes clustering at low T-scores (-3.48 average).

To maintain explainability, prototypes are periodically projected onto the nearest training examples, ensuring that each represents an actual patient case rather than abstract features.

3.4. Multi-Task Learning

Traditional approaches inadequately exploit the fact that the bone density has a continuous nature and have notable limitations. Tabular methods heavily rely on external BMD values, achieving only 59% accuracy without them (Table 3(b)). Two-stage vision frameworks that first predict BMD to make the diagnosis suffer from classification errors due to sharp clinical thresholds at $T=-1$ and $T=-2.5$ (See Table 3(c)). Direct classification ignores continuous bone-density variations, plateauing at about 76% accuracy, although not far from the standard cross-entropy approaches (See Table 3(a)).

We address these issues using multi-task learning, jointly optimising for classification and T-score regression. This

encourages the model to learn continuous bone-density representations, smooth decision boundaries, and regularise training. Formally:

$$\mathcal{L}_{ProtoMedX} = \mathcal{L}_{cls} + \lambda_1 \mathcal{L}_{reg} + \lambda_2 \mathcal{L}_{proto} \quad (2)$$

where $\mathcal{L}_{cls}$ handles classification, $\mathcal{L}_{reg}$ predicts T-scores, and $\mathcal{L}_{proto}$ is the prototype learning loss. Including regression ($\lambda_1 = 0.3$) significantly improves accuracy, removing it results in a 2.46% accuracy drop. Thus, explicitly modelling bone density continuity substantially enhances clinical predictions which, in the end, are in a form of 3-class classification (Normal/ Osteopenia/Osteoporosis). The λ multipliers were determined after grid search.

3.5. Clinical Explanations

ProtoMedX provides clinically interpretable, multi-level explanations for each diagnosis. (See Sec.5.4 and Fig. 5).

Classification Confidence: Model confidence is derived from the weighted k-NN voting among prototypes:

$$\mathcal{C} = \frac{\sum_{i \in \mathcal{N}_c} e^{-d_i/\tau}}{\sum_{j=1}^k e^{-d_j/\tau}}, \quad (3)$$

where $\mathcal{N}_c \subseteq \{1, \dots, k\}$ are the k nearest prototypes of the predicted class c , d_i is the Cosine distance to prototype i , and τ controls the sharpness of the weighting.

High confidence (e.g. $> 90\%$) indicates strong consensus; low confidence (e.g. $< 60\%$) highlights borderline or uncertain cases. An advantage of the proposed ProtoMedX is the ability to analyse the decisions.

Prototype-based Reasoning: Each decision is supported by the most similar k ($k = 3$) prototypes, displaying source patient ID, clinical features, influence weight, and prototype representativeness.

Feature-Level Analysis: For each clinical feature j , we compute relative deviation from class norms (for class c):

$$\delta_j = \frac{|x_j^{tab} - \mu_{c,j}|}{\max(\mu_{c,j}, 1)} \quad (4)$$

Features with $\delta_j > 0.5$ signal clinically atypical values.

Voting Visualisation: Bar plots show vote distributions, helping identify borderline versus confident classifications.

Misclassification Analysis: ProtoMedX provides confidence, which is an indication of a possible misclassification as described earlier. Furthermore, for the errors, in real-world settings, confidence scores and the explanations should aid the analysis. For our study, we visualized both predicted and true class prototypes, supporting error analysis. Most mistakes occur near diagnostic boundaries (e.g., Normal→Osteopenia). (See Fig. 5 for an example).

3.6. Prototype-Based Classification

During inference, ProtoMedX performs explainable classification through k-nearest neighbour search:

Algorithm 1 ProtoMedX Inference and Explanation Generation

- 1: **Input:** Patient $(\mathbf{x}^{img}, \mathbf{x}^{tab})$, Learned prototypes $\mathbf{P}$
 - 2: **// Feature extraction and fusion**
 - 3: $\mathbf{z}, \alpha \leftarrow \text{EncodeFuse}(\mathbf{x}^{img}, \mathbf{x}^{tab})$
 - 4: **// k-NN classification**
 - 5: Compute distances: $d_{c,k} = \|\mathbf{z} - \mathbf{p}_{c,k}\|_2$ for all prototypes
 - 6: Select $k=3$ nearest prototypes: $\mathcal{N} = \{(c_i, k_i, d_i)\}_{i=1}^3$
 - 7: Calculate weights: $w_i = \exp(-d_i/\tau)$, normalize to sum to 1
 - 8: Vote: $v_c = \sum_{i:c_i=c} w_i$ for each class c
 - 9: Predict: $\hat{y} = \arg \max_c v_c$
 - 10: **// Generate clinical explanation**
 - 11: Confidence: $\mathcal{C} = v_{\hat{y}}$ (winning vote weight)
 - 12: Retrieve prototype metadata: patient IDs, T-scores, clinical features
 - 13: Compute feature deviations: $\delta_j = |x_j^{tab} - \mu_{\hat{y},j}| / \max(\mu_{\hat{y},j}, 1)$
 - 14: Flag atypical features where $\delta_j > 0.5$
 - 15: **Return:** Class $\hat{y}$, confidence $\mathcal{C}$, similar patients $\mathcal{N}$,
 - 16: modality weight α , feature analysis $\{\delta_j\}$
-

This approach achieves 89.8% accuracy with k-NN classification, comparable to the less explainable approach of neural network classification layers instead of K-NN (89.6%).

4. Experiments

4.1. Dataset and Preprocessing

We utilise a clinical dataset of 4,160 lumbar spine DEXA scans with corresponding patient records. Each scan includes DEXA produced T-scores and clinical features (age, sex, height, weight, previous fractures, parent fractured hip, current smoker, glucocorticoids, rheumatoid arthritis, secondary osteoporosis, and alcohol intake).

Images undergo standardised preprocessing: adaptive histogram equalisation for contrast enhancement and Gaussian filtering ($\sigma = 1.0$) for noise reduction. All images are resized to 240×240 for CrossViT feature extraction. The dataset exhibits class imbalance typical of clinical populations: Normal (45%), Osteopenia (38%), and Osteoporosis (17%).

4.2. Training Configuration

We use 80% of the data for training (with an internal 10% validation split) and 20% as a held-out test set for evalua-

tion. Differential learning rates are set for each component: 5×10^{-5} for frozen image features, 5×10^{-4} for tabular layers trained from scratch, and 1×10^{-3} for prototype adaptation.

Training uses the AdamW optimiser with weight decay 10^{-4} , cosine annealing schedule, and early stopping (patience=15). Prototypes are initialised via k-means clustering on training features and periodically projected to the nearest training examples to maintain explainability.

4.3. Evaluation Protocol

We evaluate using: (1) overall accuracy, (2) per-class precision/recall/F1, and (3) clinical agreement (predicted vs. true diagnostic category from T-scores). For ablations, we systematically remove components and measure performance degradation.

4.4. Baseline Comparisons

We compare our method against SOTA tabular approaches, vision-only models, other prototype models (see Table 3), and recent SOTA bone health classification methods (see Table 4). For a fair comparison, we evaluated baselines and comparison in our aforementioned dataset and training splits, keeping the exact architecture, pre-processing and other architecture contributions claimed by each of the surveyed methods.

5. Results

5.1. Quantitative Performance

ProtoMedX demonstrates strong clinical alignment: 83.9% diagnostic agreement, and 91.2% sensitivity for Normal vs. Abnormal detection and SOTA performance with 89.2% accuracy on three-class bone health classification.

Detailed metrics show balanced performance: 89.3% precision, 89.2% recall, and 89.1% F1-score. Per-class analysis reveals excellent Normal detection (92.1% precision, 95.0% recall), balanced Osteopenia performance (88.7% precision, 78.6% recall), and high Osteoporosis sensitivity (79.4% precision, 92.6% recall). Crucial for identifying severe cases requiring immediate intervention.

Note that MLP-MultiTask jointly estimates T-scores and bone health classification, whereas MLP-SingleTask estimates T-scores only, which are then mapped to the classification boundaries. The multi-task approach consistently improves performance across modalities.

5.2. Ablation Study

Component-wise ablation reveals each architectural contribution:

Prototype learning and cross-modal attention provide the largest gains (4.33% and 4.09%), validating our architectural choices. Multi-task learning adds 2.46%, confirming

Table 1. Bone Health Classification Performance comparison across methods and modalities (%)

Model	Modality	Acc.	Prec.	Rec.	F1
ProtoMedX	Multi	89.8	89.3	89.2	89.1
MLP-MultiTask	Multi	87.8	84.7	85.4	85.0
MLP-SingleTask	Multi	82.7	82.1	75.0	77.6
CrossViT	Multi	76.2	75.2	64.0	68.2
MLP	Multi	72.3	65.3	66.0	65.6
MLP-MultiTask	Vision	87.8	85.1	85.4	85.2
ProtoMedX-V	Vision	87.5	83.6	85.3	84.2
MLP-SingleTask	Vision	82.3	82.1	73.3	76.1
CrossViT	Vision	74.7	69.5	67.7	68.5
XGBoost	Vision	73.2	65.7	66.0	65.8
MLP	Vision	73.2	65.8	66.1	65.9

Table 2. Ablation study of ProtoMedX components

Configuration	Accuracy (%)	Δ
Full ProtoMedX	89.8	-
w/o gate mechanism	88.72	-1.08
w/o multi-task	87.34	-2.46
w/o cross-attention	85.71	-4.09
w/o prototypes	85.47	-4.33
Baseline (no components)	84.99	-4.81

that continuous T-score supervision, alongside categorical classification, enhances overall performance.

Table 3. Empirical demonstration of key architectural choices

(a) Vision Backbone Comparison		
Backbone	Acc (%)	
ResNet/VGG	67/73	
DenseNet/ViT	71/74	
CrossViT [5]	76	
(b) Tabular Methods: BMD Dependency		
Method	w/ BMD	w/o BMD
XGBoost/GB	85	53/54
Decision Tree	83	59
Logistic Reg.	79	55
(c) BMD Regression vs Classification		
Model	MSE	Acc (%)
ViT	0.02	68.7
DenseNet	0.03	46.9
CrossViT	0.05	71.4

Validation of other architectural choices: Table 3(a) validates CrossViT selection for vision backbone. Table 3(b) reveals tabular methods’ critical dependency on BMD values, without them, accuracy drops from 85% to between 53 and 59%. This justifies our vision-centric approach, as clinical BMD measurements may be unavailable in settings where, for instance, X-Ray imagery is used instead of the more costly DEXA machines. Table 3(c) demonstrates that low BMD regression MSE does not guarantee high classification accuracy. Despite ViT achieving lowest MSE (0.02), CrossViT’s higher MSE (0.05) yields superior classification (71.4% vs 68.7%), because small BMD errors near diagnostic thresholds ($T = -1, T = -2.5$) cause misclassifications. With regression-then-prediction tasks always showing inferior classification performance. This motivates our multitask approach that jointly optimizes both objectives.

5.3. Prototype Explainability




Figure 3. T-SNE analysis of Fused Prototype feature space with 18 learned prototypes showing clear class separation.

T-SNE visualization (Figure 3) reveals well-separated prototype clusters with clear decision boundaries (Figure 4). Normal prototypes (green) cluster tightly at positive T-scores, while Osteoporosis prototypes (red) show broader distribution reflecting disease heterogeneity. The k-NN decision regions demonstrate how prototypes create explainable classification boundaries in the learned feature space (Figure 4).

5.4. Clinical Explanations

Figure 5 presents ProtoMedX’s explanation interface, which traces each prediction to the $k=3$ most similar prototypes with their clinical characteristics and similarity weights.




Figure 4. T-SNE analysis of k-NN decision boundaries ($k = 3$) in Fused Prototype feature space demonstrating how prototypes define diagnostic regions.

Panel (1) shows a correct osteoporosis classification with 91.4% confidence. All three prototypes belong to the osteoporosis class ($w = 0.523, 0.281, 0.110$), demonstrating strong consensus. The feature analysis identifies key risk factors: previous fracture ($\delta = 0.80$), glucocorticoid use ($\delta = 0.95$), and secondary osteoporosis ($\delta = 0.88$), providing clear clinical justification.

Panel (2) demonstrates the model’s transparency in error analysis. ProtoMedX incorrectly predicts osteopenia instead of osteoporosis, but the low confidence (47.0%) immediately flags uncertainty. The ambiguous voting distribution and mixed-class prototypes reveal the patient lies at a diagnostic boundary. Feature deviations ($\delta = 0.66\text{--}0.72$) are moderate, indicating weaker discriminative signals.

This framework enables clinicians to assess prediction reliability. Correct predictions average 85.3% confidence versus 48.9% for misclassifications. Most errors occur at diagnostic boundaries where patients exhibit characteristics of multiple classes, precisely where transparent decision support is most valuable.

5.5. Comparison to Competitors

ProtoMedX significantly outperforms existing approaches (Table 4), achieving 14 to 27% absolute improvement. Notably, our vision-only variant surpasses all prior methods while providing inherent explainability.

6. Conclusion

We presented ProtoMedX, the first prototype-based framework for explainable bone health classification. By rethinking machine learning approaches to bone health classification through case-based reasoning, ProtoMedX achieves 89.8% accuracy, surpassing prior methods by 14 to 27%,

1. CORRECT CLASSIFICATION

Query Patient			
	ID: 765	Age: 75 yrs	Height: 158cm
	Sex: F	Weight: 52kg	+7 others

True: Osteoporosis → Predicted: Osteoporosis
T-score: -3.1 (true) → -2.7 (predicted)

Classification Confidence

91.4%

2. MISSCLASSIFICATION

Query Patient			
	ID: 986	Age: 69 yrs	Height: 150cm
	Sex: F	Weight: 65kg	+7 others

True: Osteoporosis → Predicted: Osteopenia
T-score: -2.7 (true) → -2.4 (predicted)

Classification Confidence

47.0%

(1.a) K=3 Most Similar Prototypes

Source: 812	w=0.523
Age: 74 Sex: Female T-Score -3.1 Fracture: Yes Represents 68 similar cases	Osteoporosis

Source: 637	w=0.281
Age: 78 Sex: Female T-Score -3.5 Fracture: Yes Represents 54 similar cases	Osteoporosis

Source: 867	w=0.110
Age: 70 Sex: Female T-Score -2.8 Fracture: Yes Represents 49 similar cases	Osteoporosis

(1.b) K-NN Class Voting Distribution

N	3.2%
OP	5.4%
OS	91.4%

(1.c) Clinical Feature Analysis

Features with significant deviation ($\delta > 0.5$)

- Previous Fracture: $\delta = 0.80$
- Glucocorticoids: $\delta = 0.95$
- Secondary Osteoporosis: $\delta = 0.88$

(2.a) K=3 Most Similar Prototypes

Source: 12	w=0.614
Age: 74 Sex: Female T-Score -2.4 Fracture: Yes Represents 111 similar cases	Osteopenia

Source: 368	w=0.221
Age: 78 Sex: Female T-Score -2.7 Fracture: Yes Represents 58 similar cases	Osteoporosis

Source: 1444	w=0.165
Age: 70 Sex: Female T-Score -2.8 Fracture: Yes Represents 28 similar cases	Osteoporosis

(2.b) K-NN Class Voting Distribution

N	10.0%
OP	43.0%
OS	47.0%

(2.c) Clinical Feature Analysis

Features with significant deviation ($\delta > 0.5$)

- Previous Fracture: $\delta = 0.72$
- Glucocorticoids: $\delta = 0.67$
- Secondary Osteoporosis: $\delta = 0.66$

Figure 5. ProtoMedX Clinical Explanations. (1) Correct classification and (2) misclassification examples. Each panel includes: (a) prototype similarity with annotated clinical metadata, (b) model confidence and class voting distribution, and (c) clinical feature deviation analysis.

Table 4. Comparative analysis of Bone Health Classification approaches.

Institution	Model	Modality	Acc (%)
Ours	ProtoMedX	Multi	89.8
Ours	ProtoMedX-V	Vision	87.58
Okayama [28]	EfficientNet-b7	Vision	75.7
Ours	ProtoPNet-MT	Vision	72.4
Ours	xDNN (CrossViT backbone)	Vision	71.0
Stanford [14]	ResNet-50	Vision	69.2
Toosi [26]	ViT	Vision	68.3
Anna [32]	VGG16	Vision	66.4
Ours	ProtoPNet	Vision	64.1
Anna [32]	Inception v3	Vision	62.9

while offering inherent explainability through prototypes and decision boundaries that clinicians can interpret and critique to support their medical judgment.

Three innovations drive this performance: (1) Prototype-based reasoning which adds 4.33%, while enabling transparent explanations (2) Cross-modal attention which ensures balanced fusion, contributing 4.09%, and (3) Multi-task learning leveraging bone density continuity which improves accuracy by 2.46% over single-task classification.

Unlike *post hoc* interpretation methods, ProtoMedX is explainable-by-design through direct comparison to learned prototypes. Each prediction traces to similar prototypes with complete clinical profiles, enabling clinicians to understand the reasoning behind each classification. The dual

prototype architecture separately models visual (trabecular thinning and porous bone) and clinical features (demographics, risk factors) before fusing them for diagnosis.

Analysing the limitations of the proposed architecture, we notice that the marginal improvement from multi-modal fusion (87.58% to 89.8%) suggests unexploited complementary information. Additionally, full-image prototypes lack the granular localisation needed for precise clinical guidance. Future work would, therefore, explore explainable-by-design attention-guided prototype localisation and longitudinal modelling for monitoring disease progression.

ProtoMedX demonstrates that it is possible to achieve improved explainability without sacrificing performance. By aligning with clinical reasoning we advance toward deployable xAI that enhances, rather than replaces, clinical judgment.

7. Acknowledgments

This study is approved by the North West- Preston Research Ethics Committee (REC ref 21/NW/0309). This work is supported by ELSA – European Lighthouse on Secure and Safe AI funded by the European Union under grant agreement No. 101070617.

References

- [1] Plamen Angelov and Eduardo Soares. Towards explainable deep neural networks (xdsn). *Neural Networks*, 130:185–194, 2020. Epub 2020 Jul 11. 3
- [2] Seung Min Baik, Hye Jin Kwon, Yujin Kim, Jaehyun Lee, Young Hoon Park, and Dong Jin Park. Machine learning model for osteoporosis diagnosis based on bone turnover markers. *Health Informatics Journal*, 30(3): 14604582241270778, 2024. 3
- [3] Centers for Disease Control and Prevention, National Center for Health Statistics. National Health and Nutrition Examination Survey (NHANES). <https://www.cdc.gov/nchs/nhanes/index.html>. 2
- [4] Chaofan Chen, Oscar Li, Chaofan Tao, Alina Jade Barnett, Jonathan Su, and Cynthia Rudin. This looks like that: Deep learning for interpretable image recognition. In *Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS)*, Vancouver, Canada, 2019. 3
- [5] Chun-Fu (Richard) Chen, Quanfu Fan, and Rameswar Panda. Crossvit: Cross-attention multi-scale vision transformer for image classification. *arXiv preprint*, arXiv:2103.14899v2, 2021. 3, 6
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*, pages 1597–1607. PMLR, 2020. 4
- [7] European Union. Regulation (EU) 2024/1680 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2019/2144 and (EU) 2021/392 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:L:2024:168:TOC>, 2024. OJ L 168, 12.6.2024, p. 1–294. 1
- [8] Alicja Gosiewska and Przemyslaw Biecek. Do not trust additive explanations. *arXiv preprint arXiv:1903.11420*, 2020. 3
- [9] Danial Hooshyar and Yeongwook Yang. Problems with shap and lime in interpretable ai for education: A comparative study of post-hoc explanations and neural-symbolic rule extraction. *IEEE Access*, 2024. Accepted for publication. 3
- [10] International Osteoporosis Foundation. Epidemiology of osteoporosis and fragility fractures. <https://www.osteoporosis.foundation/facts-statistics/epidemiology-of-osteoporosis-and-fragility-fractures>. 1, 2
- [11] John A. Kanis and University of Sheffield Centre for Metabolic Bone Diseases. Frax[®] fracture risk assessment tool. <https://www.fraxplus.org/calculation-tool>. 2
- [12] Satyapriya Krishna, Tessa Han, Alex Gu, Steven Wu, Shahin Jabbari, and Himabindu Lakkaraju. The disagreement problem in explainable machine learning: A practitioner’s perspective. *arXiv preprint arXiv:2202.01602*, 2022. 3
- [13] Jae-Sung Lee, Sushil Adhikari, Li Liu, Hyun-Gi Jeong, Hae-Ryoun Kim, and Suk-Joong Yoon. Osteoporosis detection in panoramic radiographs using a deep convolutional neural network-based computer-assisted diagnosis system: a preliminary study. *Dentomaxillofacial Radiology*, 48(1): 20170344, 2019. 2
- [14] Angela Luan, Zubair Maan, Kevin Y. Lin, and Jeffrey Yao. Application of machine learning to osteoporosis and osteopenia screening using hand radiographs. *Journal of Hand Surgery (American Volume)*, 50(1):43–50, 2025. 1, 2, 3, 8
- [15] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30*, pages 4765–4774. Curran Associates, Inc., 2017. 1, 3
- [16] Hassaan Malik and Rimsha Farooq. Osteoporosis dxa scans images of spine from pakistan, 2021. 2
- [17] National Osteoporosis Guideline Group (NOGG). <https://www.nogg.org.uk>. 2
- [18] B. Ongphiphadhanakul, R. Rajatanavin, L. Chailurkit, N. Piateu, K. Teerarungsikul, R. Sirisriro, S. Komindr, and G. Pauvilai. Prediction of low bone mineral density in postmenopausal women by artificial neural network model compared to logistic regression model. *Journal of the Medical Association of Thailand*, 80(8):508–515, 1997. 3
- [19] Daria Onitii. The limits of explainability & human oversight in the eu commission’s proposal for the regulation on ai: A critical approach focusing on medical diagnostic systems. *Information & Communications Technology Law*, 32(2):170–188, 2023. 1
- [20] Tara M Pattilachan, Ugur Demir, Elif Keles, Debesh Jha, Derk Klatte, Megan Engels, Sanne Hoogenboom, Candice Bolan, Michael Wallace, and Ulas Bagci. A critical appraisal of data augmentation methods for imaging-based medical diagnosis applications. *arXiv preprint arXiv:2301.02181*, 2023. 2
- [21] Joona Pohjonen, Carolin Stürenberg, Atte Föhr, Reija Randen-Brady, Lassi Luomala, Jouni Lohi, Esa Pitkänen, Antti Rannikko, and Tuomas Mirtti. Augment like there’s no tomorrow: Consistently performing neural networks for medical imaging. *arXiv preprint arXiv:2206.15274*, 2022. 2
- [22] J. M. Queraltó, J. Torres, and M. Guinot. Neural networks for the biochemical prediction of bone mass loss. *Clinical Chemistry and Laboratory Medicine*, 37(8):831–838, 1999. 3
- [23] Royal Osteoporosis Society. Osteoporosis: Bone density scan (dxa/dexa). <https://theros.org.uk/information-and-support/osteoporosis/scans-tests-and-results/bone-density-scan-dxa/>. 2
- [24] Royal Osteoporosis Society. Written evidence submitted to the health and social care committee. <https://committees.parliament.uk/writtenevidence/112974/pdf/>, 2022. 1
- [25] 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. 1, 3

- [26] Amir Sarmadi, Zahra Sadat Razavi, Andre J van Wijnen, and Masoud Soltani. Comparative analysis of vision transformers and convolutional neural networks in osteoporosis detection from x-ray images. *Scientific Reports*, 14(1):18007, 2024. [3](#), [8](#)
- [27] 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. *arXiv preprint arXiv:1610.02391*, 2019. [1](#), [3](#)
- [28] Shintaro Sukegawa, Ai Fujimura, Akira Taguchi, Norio Yamamoto, Akira Kitamura, Ryosuke Goto, Keisuke Nakano, Kiyofumi Takabatake, Hotaka Kawai, Hitoshi Nagatsuka, and Yoshihiko Furuki. Identification of osteoporosis using ensemble deep learning model with panoramic radiographs and clinical covariates. *Scientific Reports*, 12(1):6088, 2022. [1](#), [3](#), [8](#)
- [29] Ronan Tegerdine, Frankie Willard, Luke Triplett, Dennis Tang, Luke Moffett, Jon Donnelly, Alina Jade Barnett, and Cynthia Rudin. Protopnet and descendants, 2023. [3](#)
- [30] S. Tomasevic-Todorovic, A. Vazic, A. Issaka, and F. Hanna. Comparative assessment of fracture risk among osteoporosis and osteopenia patients: a cross-sectional study. *Open Access Rheumatology: Research and Reviews*, 10:61–66, 2018. [1](#)
- [31] Adam Tupper and Christian Gagné. Analyzing data augmentation for medical images: A case study in ultrasound images. *arXiv preprint arXiv:2403.09828*, 2024. [2](#)
- [32] P Varalakshmi, Smruthi Sathyamoorthy, V Darshan, V Ramanujan, and Sakthi Jaya Sundar Rajasekar. Detection of osteoporosis with dexa scan images using deep learning models. In *2022 International Conference on Advances in Computing, Communication and Applied Informatics (ACCAI)*, pages 1–6, 2022. [1](#), [2](#), [3](#), [8](#)
- [33] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *European Conference on Computer Vision (ECCV)*, pages 499–515. Springer, 2016. ECCV 2016. [4](#)
- [34] Mardi Witzel, Gaston H. Gonnet, and Tim Snider. Explainable ai policy: It is time to challenge post hoc explanations. CIGI Papers No. 296, Centre for International Governance Innovation (CIGI), 2024. Licensed under CC BY 4.0. [3](#)
- [35] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European Conference on Computer Vision*, pages 818–833. Springer, 2014. [3](#)
- [36] Bin Zhang, Keyan Yu, Zhenyuan Ning, Ke Wang, Yuhao Dong, Xian Liu, Shuxue Liu, Jian Wang, Cuiling Zhu, Qinqin Yu, Yuwen Duan, Siying Lv, Xintao Zhang, Yanjun Chen, Xiaojia Wang, Jie Shen, Jia Peng, Qiuying Chen, Yu Zhang, Xiaodong Zhang, and Shuixing Zhang. Deep learning of lumbar spine x-ray for osteopenia and osteoporosis screening: A multicenter retrospective cohort study. *Bone*, 140:115561, 2020. [1](#), [2](#), [3](#)
- [37] Kang Zheng, Yirui Wang, Xiao-Yun Zhou, Fakai Wang, Le Lu, Chihung Lin, Lingyun Huang, Guotong Xie, Jing Xiao, Chang-Fu Kuo, and Shun Miao. Semi-supervised learning for bone mineral density estimation in hip x-ray images. In

International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI), pages 152–162. Springer, 2021. [4](#)