# Adaptive Neural Ensemble for X-Ray Preprocessing in Clinical Decision Support Systems

I.I. Mishchenko (0009-0007-5328-7729) · A.E. Misnik (0000-0003-4750-8138)
Belarusian-Russian University, Mogilev, Belarus

# Problem & Motivation

- AI diagnosis quality follows «garbage in – garbage out»
- X-rays suffer from noise, artifacts, non-uniform exposure
- Hip-Spine syndrome requires precise landmarks on paired views

Static pipelines fail twice:

1) Degraded scans — insufficient correction

2) Good scans — over-filtering erases subtle pathological cues

Our idea

DICOM metadata (kVp, mAs, distance) encodes expected noise and exposure — but is rarely used.

Route each image through only the correction modules it actually needs.

# Problem & Motivation

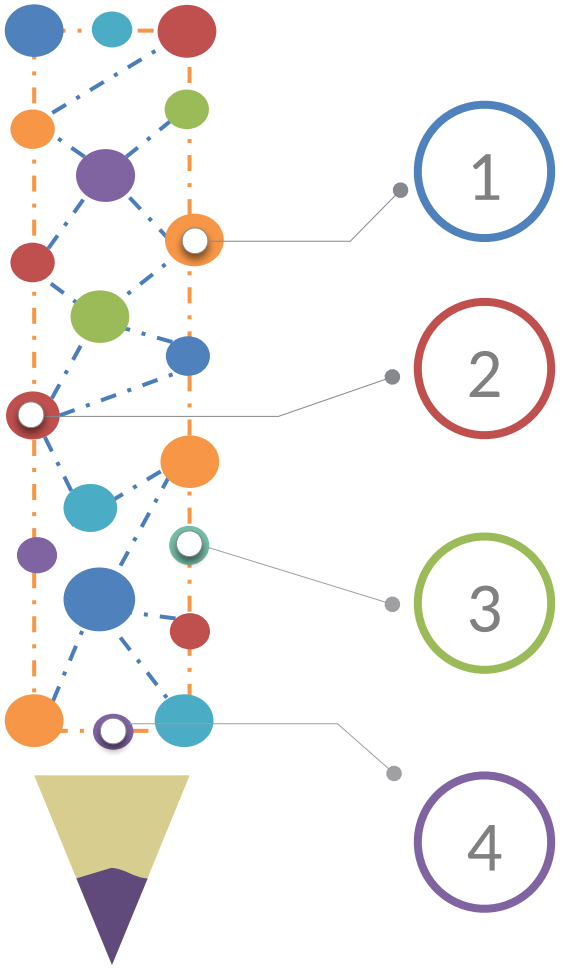
**Typical defects:**

- **Noise (pulse, Gaussian).**
- **Low contrast.**
- **Uneven brightness.**
- **Artifacts.**




**Selecting methods and their parameters manually is a complex and labor-intensive task that depends on the specific image.**

# Proposed Architecture




## Image preprocessing

DICOM → PNG conversion, noise filtering, CLAHE, normalization

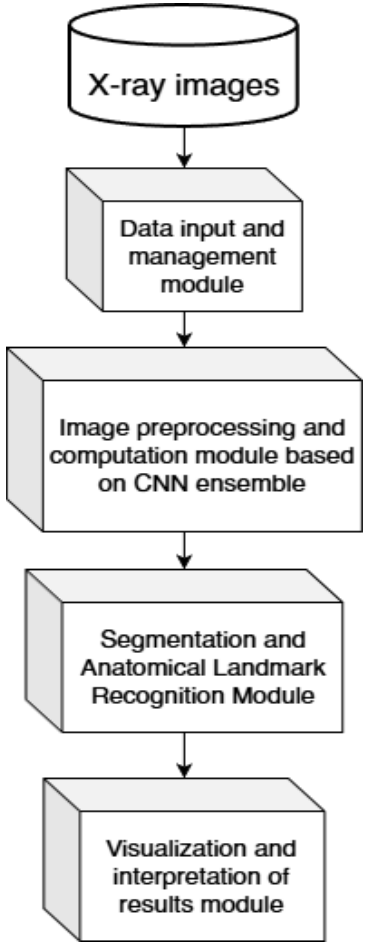
## Segmentation and analysis (YOLOv11)

Detection of the spine and vertebrae

## Automatic diagnosis

Calculation of angular deviations
(scoliosis, kyphosis, Hip-spine syndrome)

## Visualization and storage of data



# Proposed Architecture: Supervisor + Specialized Modules

**Supervisor Network**

ResNet-50 (depthwise separable) + cross-attention fusion of DICOM metadata (LSTM) and image statistics

→

**Dynamic Routing**

Deep Q-Network chooses activation mask — only needed modules run

→

**Specialized Modules**

Contrast • Brightness • Denoising • Geometric • Scaling

→

**QoI Early Termination**

Skip remaining stages when quality is already sufficient

Synthetic data generation: Poisson + Gaussian noise and metallic artifacts injected into clean radiographs to train the denoising autoencoder robustly.

# Quality-of-Image Metric & Early Termination

$$QoI = 0.25 \cdot NIQE + 0.25 \cdot BRISQUE + 0.30 \cdot Edge\ Coherence + 0.20 \cdot SNR$$

- Weights tuned by grid search
- Correlation with radiologist ratings: r = 0.87
- Early exit if QoI gain < 0.03 and QoI > 0.92

Efficiency gains

- −15–20% average inference time
- −50% FLOPs (4.2 → 2.1 GFLOPs)
- Ablation: disabling early exit adds only +0.5% accuracy but +45% time
- Suitable for edge / IoT medical deployment

# Results: Segmentation (6,317 X-rays, 10-fold CV)

| Method | Landmark Acc. | Recall | Dice | IoU |
|---|---|---|---|---|
| No preprocessing | 68% | 59% | 71% | 64% |
| CLAHE + filters | 81% | 77% | 83% | 76% |
| **Proposed ensemble** | **93%** | **90%** | **91%** | **88%** |

**+12–15% over CLAHE, all improvements significant ($p < 0.01$)**

External validation (MURA, n = 892):

91.2% vs 82.8% baseline ($p < 0.05$)

Confirms cross-institutional robustness

# Results: Classification & Ablations

| Method | Accuracy | F1 |
|---|---|---|
| No preprocessing | 72% | 70% |
| CLAHE + filters | 84% | 82% |
| **Proposed ensemble** | **93%** | **92%** |

Error analysis: CLAHE suppresses low-contrast osteophytes and creates halo artifacts near implants; the ensemble reduces implant artifacts by 25%.

Ablation study

- w/o DICOM metadata: −4.2% landmark accuracy → metadata routing is essential
- w/o early termination: 348 ms vs 240 ms, accuracy +0.5% only
- Failure mode (3.2% of scans): extreme metallic scatter misrouted — future MAR submodule

# Conclusion & Future Work

Contributions

1. Supervisor-driven ensemble with metadata-aware routing
2. Formalized QoI metric + early termination (−50% FLOPs)
3. +12–15% segmentation, 93% classification accuracy

Future work

- Clinical deployment: Docker/Kubernetes microservice via DICOMweb into PACS
- Federated Learning for privacy-preserving multi-center updates
- CT/MRI cross-modality adaptation and 3D reconstruction

Thank you for your attention!