





01 The Clinical Bottleneck

ANATOMICAL LANDMARK • DATA-SCARCE • CLINICAL TRUST

Reliable anatomical landmark detection is **not just about accuracy**, it requires representations that remain **robust and uncertainty-aware** even with only a handful of annotations.

02 The CDPM-Align Framework

988 X-RAYS • 3 ANATOMICAL DISTRICTS • MULTI-SCALE




- Generative Pretraining:** A **conditional diffusion** model is trained on small-scale heterogeneous datasets to learn rich anatomical priors without any landmark annotations.
- Multi-Scale Guidance Alignment:** Enforces feature consistency across U-Net hierarchies and diffusion timesteps, using **classifier-free guidance** as a structural descriptor to learn stable spatial representations.
- Few-Shot Downstream Fine-Tuning:** Fine-tunes the pretrained encoder with **minimal annotations** for accurate, robust, and uncertainty-calibrated landmark detection.

METRICS:

- MRE $\rightarrow$ Accuracy** (Mean radial error)
- ERE $\rightarrow$ Uncertainty** (Expected spatial error)
- P95 $\rightarrow$ Robustness** (95th percentile error)

03 Robustness in Low-Data Regimes

10-SHOT • 5 INDEPENDENT RUNS • $P \leq 0.001$



04 The Bottom Line

FEW-SHOT • TIGHT BOUNDS • ROBUST FEATURES

Through multi-scale guidance alignment, **CDPM-align learns robust anatomical representations** that enable **accurate and well-calibrated landmark detection** in extremely low-data regimes.