Uncertainty-Aware Foundation Models for Trustworthy Chest X-ray Report Generation

L. Julián Lechuga López [1,2] and Farah E. Shamout [1,2]

[1] New York University, [2] New York University Abu Dhabi

Main Takeaways

- Most AI-driven systems in healthcare today, rely on single-modality inputs, **without expressing uncertainty, confidence,** or diagnostic risk.
- Our goal is to develop multimodal models that generate interpretable reports and estimate the relative contribution and uncertainty of each input modality, emphasizing **explainability and clinical utility**.
- We envision a new evaluation strategy combining **uncertainty**, **clinical semantics**, and **expert-derived criteria** to go beyond AUC, BLEU, or ROUGE.
- We define a **chain-of-thought system** for differential diagnostics based on confidence and clinical rationale.

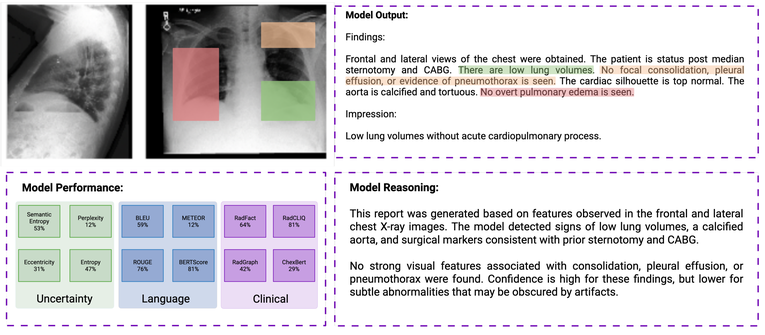

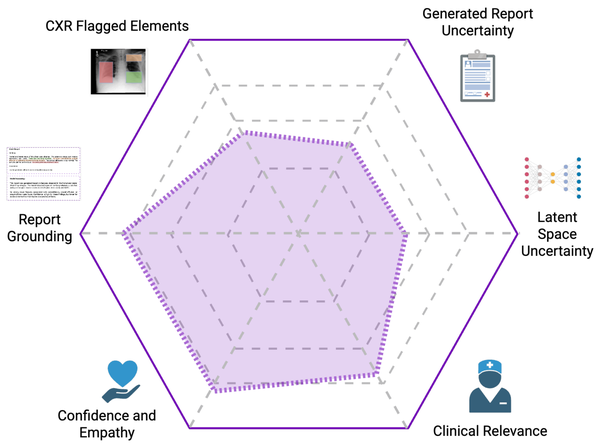

Figure 1: Interpretable Uncertainty-Aware and Clinical Evaluation Framework.

Uncertainty–Aware Report Generation

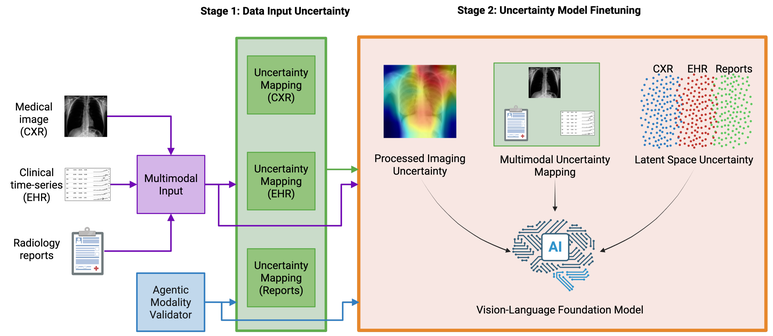

Figure 2: Framework for Uncertainty Attribution in Multimodal Clinical AI

- **Multimodal Input Attribution:** We develop a framework that disentangles and attributes uncertainty across available modalities (e.g., CXR, EHR, reports), enabling more targeted and interpretable diagnostic insights.
- **Fine-Grained Uncertainty Decision-Making:** The model captures uncertainty from the data inputs, across the generative process, and at the output prediction level, supporting more nuanced clinical reasoning.
- **Agentic Validator Integration for Clinical Oversight:** We incorporate external medical knowledge sources through validator agents to independently assess prediction plausibility and enhance clinical trust and safety.
- **Cross-Modality Uncertainty Mapping and Propagation:** Our system tracks how uncertainty flows across modalities and model stages, surfacing early indicators of diagnostic ambiguity and guiding downstream decisions.
- **System Modularity Across Clinical Tasks:** Our goal is to develop a system that can be easily adapted for use in different report generation tasks (e.g., breast cancer) and clinical data modalities.

From Prediction to Planning

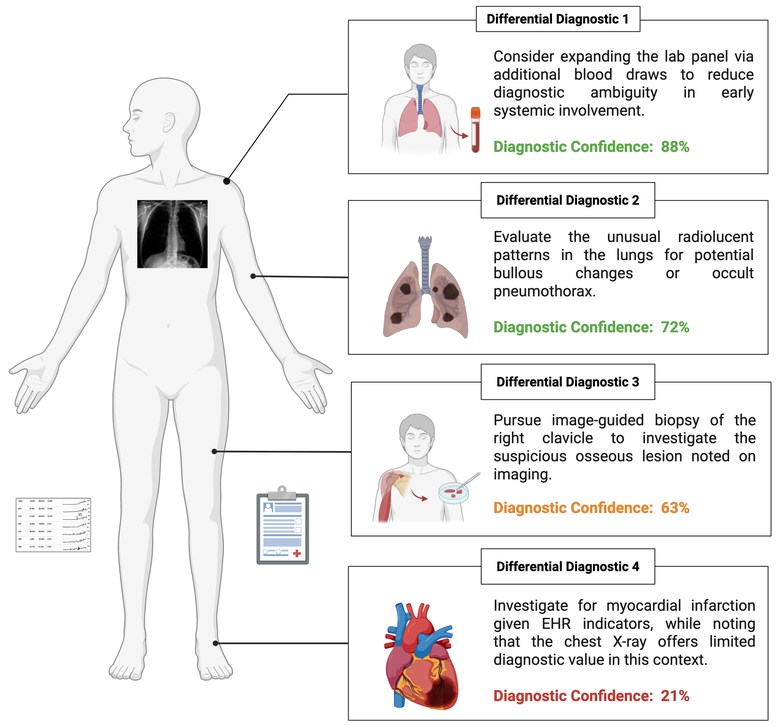

Figure 3: Simulated Uncertainty-Guided Planning.

- **Confidence-Ranked Differential Diagnoses:** The model proposes multiple diagnostic hypotheses prioritized by confidence scores, grounded in multimodal and clinical evidence, and uncertainty quantification.
- **Forward Diagnostic Reasoning and Next Steps:** Each prediction includes rationale and recommended actions (e.g., labs, imaging) to reduce uncertainty and guide clinical decision-making.
- **Interactive and Deployable in Low-Resource Settings:** Designed for human-in-the-loop use with explainability and prioritization at its core, especially under data sparsity or limited resources.