

000 001 002 003 004 005 LUNGUAGE: A BENCHMARK FOR STRUCTURED AND 006 SEQUENTIAL CHEST X-RAY INTERPRETATION 007 008 009

010 **Anonymous authors**
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053

Paper under double-blind review

ABSTRACT

Radiology reports convey detailed clinical observations and capture diagnostic reasoning that evolves over time. However, existing evaluation methods are limited to single-report settings and rely on coarse metrics that fail to capture fine-grained clinical semantics and temporal dependencies. We introduce LUNGUAGE, a benchmark dataset for structured radiology report generation that supports both single-report evaluation and longitudinal patient-level assessment across multiple studies. It contains 1,473 annotated chest X-ray reports, each reviewed by experts, and 186 of them contain longitudinal annotations to capture disease progression and inter-study intervals, also reviewed by experts. Using this benchmark, we develop a two-stage structuring framework that transforms generated reports into fine-grained, schema-aligned structured reports, enabling longitudinal interpretation. We also propose LUNGUAGESCORE, an interpretable metric that compares structured outputs at the entity, relation, and attribute level while modeling temporal consistency across patient timelines. These contributions establish the first benchmark dataset, structuring framework, and evaluation metric for sequential radiology reporting, with empirical results demonstrating that LUNGUAGESCORE effectively supports structured report evaluation. Code and data are available at: <https://anonymous.4open.science/r/lunguage>

1 INTRODUCTION

Radiology reports play a critical role in diagnosis by recording patient history, describing imaging findings, documenting procedures, and noting temporal changes. However, because they are written in unstructured free text, reports vary widely in terminology, style, and level of detail across radiologists, complicating consistent computational interpretation and hindering automated systems for report generation and evaluation. To address these challenges, structuring frameworks have been developed to convert free-text reports into standardized, machine-friendly formats (Jain et al. (2021); Khanna et al. (2023); Wu et al. (2021); Zhang et al. (2023); Zhao et al. (2024)). While these frameworks improve representational consistency, current evaluation methods remain fundamentally limited in two key aspects: temporal reasoning and fine-grained clinical accuracy.

Temporal reasoning is central to radiologic interpretation, as diagnoses often depend on comparing current and prior studies to assess whether a finding has progressed. However, most evaluation protocols (Bannur et al. (2024); Huang et al. (2024); Jain et al. (2021); Khanna et al. (2023); Ostmeier et al. (2024); Smit et al. (2020); Wu et al. (2021); Yu et al. (2023a); Zhang et al. (2023); Zhao et al. (2024)) assess reports in isolation, without incorporating previous findings. This makes it impossible to determine whether temporal expressions—such as “no change,” “improved,” or “new”—are appropriate. For instance, the statement “no change in pneumonia” cannot be meaningfully evaluated without confirming whether pneumonia was present in prior studies.

Fine-grained clinical accuracy is equally critical. Reliable interpretation depends on attributes such as precise location (e.g., “carina above 3 cm”) and lesion size (e.g., “2.5 cm”). These details are essential for diagnostic specificity and downstream decision-making, yet most evaluation protocols collapse them into broad categories. For instance, “2.5 cm right upper lobe nodule with spiculated margins” may be reduced to simply “nodule,” and this loss of granularity makes it difficult to distinguish precise from incomplete outputs.

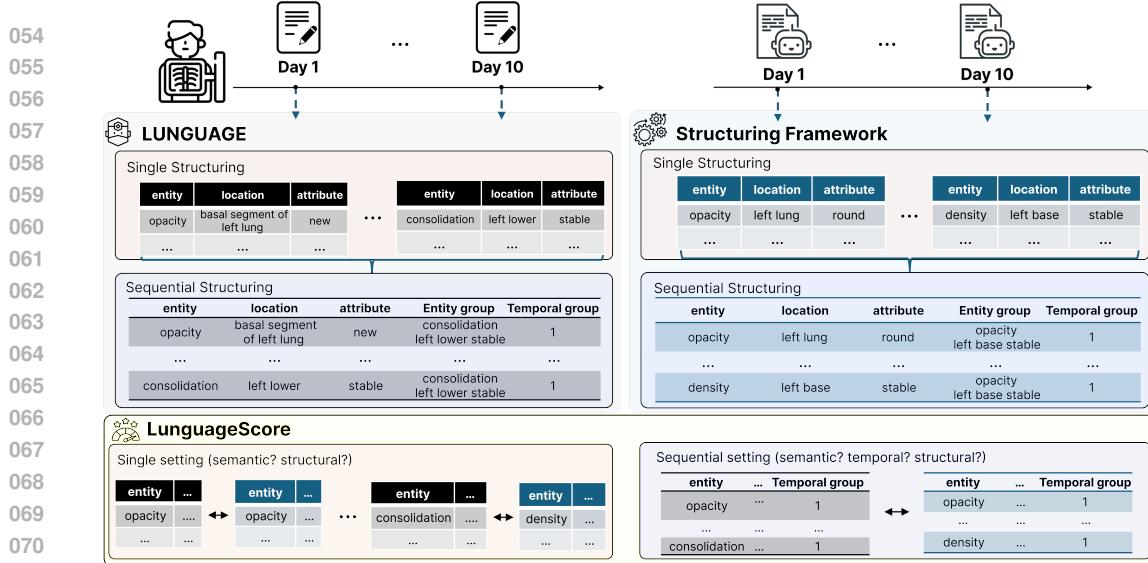


Figure 1: **Evaluation pipeline for radiology report generation.** We introduce the first evaluation framework for radiology report generation, enabling both detailed single-report assessment and comprehensive patient-level trajectory evaluation. On the left, we release **LUNGUAGE**, a radiologist-annotated benchmark of structured single and sequential chest X-ray reports. On the right, we develop a two-stage **structuring framework** that converts free-text into schema-aligned structures at both single and sequential levels. At the bottom, we present **LUNGUESCORE**, a clinically validated metric that jointly measures semantic accuracy, structural fidelity, and temporal alignment, providing clinically faithful evaluation.

Structuring frameworks have attempted to address these issues by extracting entities and relations from reports (Jain et al. (2021); Khanna et al. (2023); Wu et al. (2021); Zhang et al. (2023); Zhao et al. (2024)). Some extend this by tagging temporal descriptors such as “worsened” or “stable” (Khanna et al. (2023); Wu et al. (2021)). Yet, they remain restricted to single reports and rely only on explicitly stated expressions, without verifying consistency across time. Consequently, they cannot ensure whether findings align with prior studies or capture coherent clinical trajectories, and often miss the clinical granularity needed for precise diagnostic interpretation.

Recent report generation models have begun incorporating temporal inputs such as prior reports, imaging, or clinical indications (Bannur et al. (2024); Zhou et al. (2024)), enabling outputs that are more context-aware and temporally coherent. However, evaluation protocols have not kept pace. Generated reports are still judged at isolated timepoints rather than across a continuous timeline, making it impossible to assess whether models appropriately incorporated prior findings or preserved clinically important details at both temporal and semantic levels.

To address these limitations, we present the first evaluation pipeline for assessing radiology report generation in both single and sequential settings. Our contributions are threefold. (1) We construct **LUNGUAGE**, a fine-grained benchmark that establishes reliable ground truth for evaluation. It consists of **1,473 single reports** from 230 patients (annotated with 17,949 entities and 23,307 relation–attribute pairs across 18 clinically grounded relation types) and **186 sequential reports** from 30 patients (95,404 observation pairs across 2–14 reports per patient). These support longitudinal analysis through **ENTITYGROUPS** (linking the same finding across reports) and **TEMPORALGROUPS** (segmenting diagnostic episodes). (2) To enable automatic benchmarking on this scale, we develop a **structuring framework** that converts free text into *entity–relation–attribute* triplets and links them across time following the LUNGUAGE schema. It achieves high agreement with expert annotations (F1: 0.94 for entity–relation, 0.86 for full triplets, 0.69 for ENTITYGROUP, 0.87 for TEMPORALGROUP). (3) Building on this foundation, we introduce **LUNGUESCORE**, a clinically grounded metric that compares structured representations from generated and reference reports. Unlike prior approaches, this metric simultaneously captures semantic accuracy, structural fidelity, and temporal coherence. To our best knowledge, this is the first study to combine the highest schema granularity with explicit modeling of full diagnostic trajectories.

108

2 RELATED WORK

110 **Structuring Radiology Reports** Radiology reports encode layered clinical semantics, spanning
 111 history, imaging observations, and diagnostic reasoning. Rule-based systems (Wu et al. (2021); Zhang
 112 et al. (2023)) achieve high precision in constrained settings but struggle to generalize due to linguistic
 113 variability. Supervised transformer-based methods (Jain et al. (2021); Khanna et al. (2023); Zhao
 114 et al. (2024)) are more flexible but depend heavily on the coverage and granularity of their annotation
 115 schema. Recently, prompting-based approaches have leveraged large language models (LLMs), such
 116 as GPT-4 (Achiam et al. (2023)) and open-source variants (Liu et al. (2024); Touvron et al. (2023)), to
 117 directly produce structured outputs from free text (Busch et al. (2024); Dorfner et al. (2024); Hartsock
 118 et al. (2025); Woźnicki et al. (2024)). While these models demonstrate strong few-shot performance,
 119 they remain prone to hallucinations, inconsistent terminology, and prompt sensitivity. To mitigate
 120 this, we employ a task-specific vocabulary and schema-aligned reference set, constraining outputs to
 121 valid clinical concepts and enhancing consistency through retrieval-augmented prompting. A detailed
 122 comparison is provided in Appendix A.4.

123 **Evaluation Metrics for Radiology Report Understanding** Existing metrics fall into three cate-
 124 gories: lexical, model-based, and structure-based. Lexical metrics (BLEU (Papineni et al. (2002)),
 125 ROUGE (Lin (2004)), METEOR (Banerjee & Lavie (2005))) rely on surface overlap and often
 126 miss clinical meaning. Model-based metrics (CheXbert (Smit et al. (2020)), BERTScore (Zhang
 127 et al. (2019))) capture semantic similarity but overlook fine-grained detail. Structure-based metrics
 128 (RadGraph-F1 (Jain et al. (2021)), RaTEScore (Zhao et al. (2024))) add granularity by matching
 129 entities and relations. Recent efforts emphasize clinical error detection: ReXVal (Yu et al. (2023b))
 130 introduced expert-labeled errors, informing RadCliQ (Yu et al. (2023a)), which combines BERTScore
 131 and RadGraph-F1, while LLM-based metrics (GREEN (Ostmeier et al. (2024)), FineRadScore
 132 (Huang et al. (2024)), RadFact (Bannur et al. (2024)), CheXprompt (Zambrano Chaves et al. (2025)))
 133 aim to approximate expert judgment or factual correctness. However, most metrics still evaluate
 134 reports in isolation, overlooking temporal consistency across studies and neglecting attributes like
 135 location, extent, or progression. In contrast, our evaluation pipeline provides structured, tempo-
 136 rally aligned evaluation over patient report sequences, enabling clinically faithful assessment across
 137 semantic, structural, and temporal dimensions.

138

3 LUNGUAGE: SINGLE AND SEQUENTIAL STRUCTURED REPORTS

140 Radiology reports vary in depth and nuance, with differences in phrasing, certainty, and cross-
 141 sentence connections that make structured interpretation challenging. They are commonly divided
 142 into *indication/history*, which provides contextual cues (e.g., “history of cough”), and *findings* and
 143 *impression*, which contain detailed descriptions and diagnostic reasoning (e.g., “left opacities likely
 144 consolidation or pneumonia”). To address this complexity, we present LUNGUAGE, a benchmark
 145 dataset of radiologist-annotated chest X-ray reports in two complementary versions: 1,473 single
 146 reports and 186 sequential reports. Reports were structured through a rigorous annotation process
 147 (Appendix A.3) guided by three principles: diagnostic source distinction (separating image-based
 148 from context-based findings), semantic precision (capturing descriptive cues such as certainty, status,
 149 and other fine-grained attributes including location, severity, and morphology), and longitudinal
 150 linkage (capturing temporal consistency through entity and temporal grouping). This design reflects
 151 physicians’ clinical perspectives and supports both a single-report schema for fine-grained interpreta-
 152 tion and a sequential schema for modeling patient-level diagnostic trajectories. Figure 2 illustrates
 153 these schemas.

154

3.1 SINGLE STRUCTURED REPORT: FINE-GRAINED SCHEMA AND ANNOTATION

155 We propose a single-report schema that captures the internal structure of radiology reports by struc-
 156 turing clinically relevant information into two units: **entities**, representing core clinical concepts, and
 157 **relations**, encoding their attributes and interconnections, enabling systematic modeling of both de-
 158 tailed descriptions and cross-sentence reasoning for fine-grained and clinically faithful interpretation.

159 **ENTITIES** are assigned to one of six clinically grounded categories based on their derivability from
 160 chest X-ray imaging: PF (PERCEPTUAL FINDINGS) for directly observable image features (e.g.,

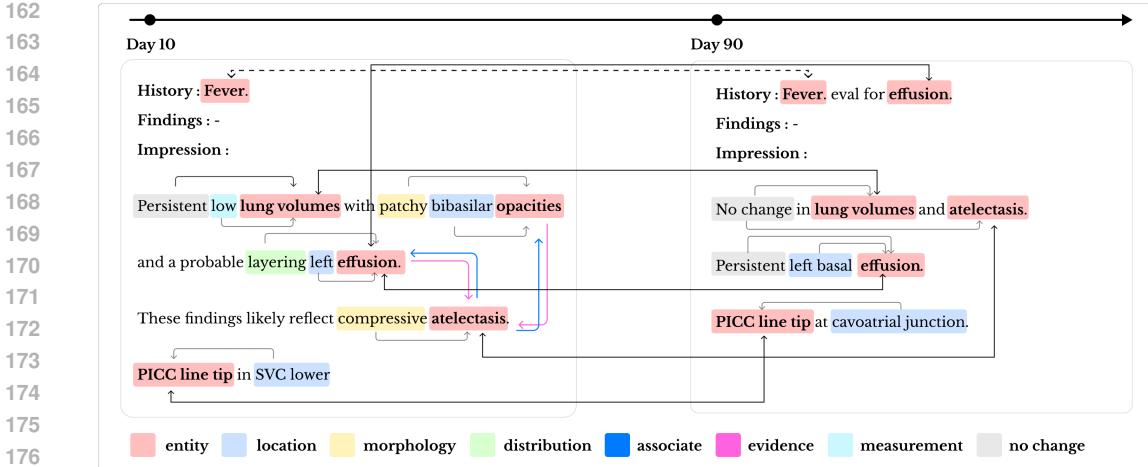


Figure 2: **Schema for Single and Sequential Report Structuring.** The figure shows two reports from the same patient at day 10 and day 90. For the single report schema (within each report), gray solid lines connect entities to attributes, while pink and blue solid lines represent inter-entity reasoning relations (ASSOCIATE, EVIDENCE). For the sequential schema (across reports), black solid lines denote entities in the same ENTITYGROUP (same clinical finding over time) and TEMPORALGROUP (same diagnostic episodes), while black dashed lines show entities in the same ENTITYGROUP but different TEMPORALGROUPS (different diagnostic episodes).

“lung,” “opacity”); CF (CONTEXTUAL FINDINGS) for diagnoses inferred from external clinical context (e.g., “pneumonia”); OTH (OTHER OBJECTS) for mentioned devices or procedures (e.g., “ET tube”); COF (CLINICAL OBJECTIVE FINDINGS) for structured observations from non-imaging sources (e.g., lab tests); NCD (NON-CXR DIAGNOSIS) for diagnoses based on other modalities (e.g., “AIDS”); and PATIENT INFO for reported history or symptoms (e.g., “fever,” “cough”).

RELATIONS capture clinical properties and inter-entity connections, often spanning multiple sentences. The schema includes diagnostic stance (DXSTATUS, DxCERTAINTY); spatial and descriptive characteristics (LOCATION, MORPHOLOGY, DISTRIBUTION, MEASUREMENT, SEVERITY, COMPARISON); temporal dynamics (ONSET, IMPROVED, WORSENED, NOCHANGE, PLACEMENT); and contextual information (PASTHx, OTHERSOURCE, ASSESSMENTLIMITATIONS)¹. It also includes two reasoning relations: ASSOCIATE (bidirectional links between related entities) and EVIDENCE (asymmetric support from a finding to a diagnosis). For example, in “left lung opacity suggests pneumonia,” the schema identifies both ASSOCIATE between *opacity* and *pneumonia*, and EVIDENCE indicating that *pneumonia* is inferred from *opacity*. Full definitions can be found in Appendix A.1.

Single Report Annotation Process We developed a two-stage annotation pipeline for 1,473 radiology reports from MIMIC-CXR (Johnson et al. (2019)) to ensure fine-grained and clinically grounded structuring of report language. The process began with schema design and initial structured drafts generated by GPT-4 (0613)², from which all candidate entity and relation terms were collected to build a comprehensive vocabulary. Four radiologists independently reviewed these terms in a blinded manner, resolving discrepancies through majority voting and consensus meetings, and referring to the Fleischner Society (Bankier et al. (2024)) terminology and UMLS (Bodenreider (2004)) mappings where appropriate. This stage unified terminology and eliminated category-level inconsistencies in advance. In the second stage, annotators revised all reports using this curated vocabulary, focusing on contextual interpretation and correction of potential LLM errors rather than category disputes. Reports were evenly divided among radiologists, who verified every (*entity, relation, attribute*) triplet, including cross-sentence relations such as ASSOCIATE and EVIDENCE. This two-step process yielded 17,949 entity instances and 23,307 relation instances, providing a

¹**Abbreviations:** “Dx” stands for “diagnosis” and is used in relations such as DXSTATUS (i.e., positive or negative finding) and DxCERTAINTY (i.e., definitive or tentative). “Hx” in PASTHx stands for “history”.

²All large language model (LLM) usage, including GPT-4, was conducted using HIPAA-compliant deployments provided by Azure and Fireworks AI.

216 reliable and clinically validated dataset for benchmarking structured report interpretation. Details of
 217 the vocabulary and annotation process are provided in Appendix A.2 and A.3.1.
 218

219 3.2 SEQUENTIAL STRUCTURED REPORT: DISEASE TRAJECTORY SCHEMA AND ANNOTATION

220 Longitudinal reports often exhibit lexical variation, abstraction shifts, and inconsistent phrasing
 221 (Meystre et al. (2008); Wang et al. (2018)). The same pathology may be described differently across
 222 timepoints, such as “left opacity” and “left lower consolidation,” with differences in wording and
 223 specificity that complicate semantic alignment and temporal reasoning. To address this, we introduce
 224 a schema that structures reports across patient timelines through two key components:
 225

226 **ENTITYGROUPS** identify observations that refer to the same underlying clinical finding, even when
 227 expressed using different terms, anatomical references, or levels of abstraction. Within each patient,
 228 all observation pairs are compared to detect semantic equivalence, regardless of when they appear
 229 in the timeline, whether the finding is reported as present or absent (DXSTATUS), or whether it is
 230 stated definitively or tentatively (DxCertainty). For example, “PICC line tip in lower SVC” and
 231 “at the cavoatrial junction” (Figure 2) may describe the same catheter tip location, reflecting inherent
 232 ambiguity in 2D imaging. Similarly, “lung volumes” reported as low on day 10 and described as “no
 233 change” on day 90 can be grouped to indicate persistent low lung volume.

234 **TEMPORALGROUPS** divide each ENTITYGROUP into distinct diagnostic episodes based on temporal
 235 distance, status shifts or certainty, and clinical change expressions (e.g., “worsening”). This approach
 236 captures clinically meaningful transitions in a patient’s condition (Chapman et al. (2011); Savova et al.
 237 (2010)). For example, “fever” mentioned in both the day 10 and day 90 reports (Figure 2) appears in
 238 the “history” section but occurs far apart in time, so treating them as separate temporal groups better
 239 reflects clinical reasoning, whereas repeated descriptions of an effusion would remain in the same
 240 group. Together, these components support fine-grained evaluation of both semantic consistency and
 241 temporal coherence in longitudinal model outputs.

242 **Sequential Report Annotation Process** We annotated 186 chest X-ray reports from 30 patients—a
 243 subset of the 230-patient cohort used in single-report annotation—to construct a gold-standard dataset
 244 for patient-level longitudinal evaluation. The same four radiologists independently reviewed reports
 245 in chronological order, linking observations that referred to the same underlying finding into ENTITY-
 246 GROUPS (e.g., “pleural effusion right lung increasing”) and dividing them into TEMPORALGROUPS
 247 (labeled 1, 2, . . .) to distinguish diagnostic episodes. Terminology was normalized when appropriate
 248 (e.g., aligning “clavicle hardware” with “orthopedic side plate”) while preserving abstraction and
 249 anatomical distinctions. Patients contributed 2–14 reports, with intervals spanning 1–1,200 days.
 250 For each patient, all observation pairs (29–141 per case) were compared, yielding 95,404 total
 251 comparisons. This rigorous process ensured both longitudinal consistency and clinically meaningful
 252 transitions such as resolution or recurrence. Further details are provided in Appendix A.3.2.

253 4 STRUCTURING FRAMEWORK FOR SINGLE AND SEQUENTIAL REPORTS

254 We develop a two-stage structuring framework that automatically structures radiology reports using the
 255 same schema as LANGUAGE, generating radiologist-like structured outputs for consistent evaluation.
 256 The framework covers both single-report and longitudinal settings, producing representations for
 257 semantic, structural, and temporal evaluation (Figure 1).
 258

259 **(i) Single Structuring Framework** To generate structured outputs from free text, we apply corpus-
 260 guided relation extraction with a LLM, which extracts (*entity, relation, attribute*) triplets aligned to
 261 our schema. The task requires handling both intra- and inter-sentential contexts and accommodating
 262 lexical variation without relying on templates. While LLMs can capture diverse phrasing and nuanced
 263 expressions, they are prone to hallucinations and inconsistencies (Busch et al. (2024); Dorfner et al.
 264 (2024); Hartsock et al. (2025); Woźnicki et al. (2024)). To mitigate errors, we guide the model with a
 265 curated vocabulary derived from our annotation corpus (Section 3.1). Details of the prompts and the
 266 vocabulary-matching algorithm are provided in Appendix B.1 and B.1.1.

267 **(ii) Sequential Structuring Framework** Building on the outputs from stage (i), we use the LLM
 268 to interpret report sequences over time. To address longitudinal variability, the model performs
 269 normalization and temporal aggregation. Each entity and its attributes are linearized into flattened text
 in chronological order relative to the initial study (e.g., “day 0: opacity right lung,” “day 30: opacity

right basilar”). The LLM is guided with few-shot examples illustrating lexical variation, abstraction shifts (e.g., descriptive to diagnostic terms), and rephrasings of persistent devices. Using these, it determines whether observations across time represent the same underlying finding and whether they belong to a single temporal group. Decisions are guided by semantic similarity, anatomical alignment, and temporal continuity. Observations reflecting recurrence after resolution or clinically disconnected events are treated as distinct temporal groups. This process yields two outputs: ENTITY GROUPS and TEMPORAL GROUPS, consistent with Section 3.2. The format combines entity, location, and temporal pattern (e.g., “pleural effusion right lung no change”), with groups numbered sequentially (1, 2, 3, . . .). This framework enables faithful structuring of longitudinal narratives, capturing clinically meaningful trajectories across report sequences. Full prompt examples are provided in Appendix B.2.

5 LUNGUESCORE: A FINE-GRAINED PATIENT-LEVEL METRIC

We propose **LUNGUESCORE**, a fine-grained metric for quantifying radiology report quality across semantic equivalence, temporal coherence, and attribute-level similarity. It captures clinically meaningful distinctions in terminology (e.g., “right clavicle hardware” vs. “orthopedic side plate”), longitudinal trends (e.g., resolution vs. decrease), and detailed attributes (e.g., 2.3 cm vs. 3.0 cm). These dimensions are integrated into a single similarity score that compares candidate and reference reports—either individually or as sequences—enabling patient-level evaluation.

Evaluation Principles. LUNGUESCORE is grounded in three clinical principles: **semantic sensitivity** captures concept-level equivalence across linguistic variation (Meystre et al. (2008); Wang et al. (2018)); **temporal coherence** ensures alignment with clinical timelines for assessing disease progression (Chapman et al. (2011); Savova et al. (2010)); and **structural granularity** evaluates fine-grained attributes critical for diagnosis (Demner-Fushman et al. (2009); Pons et al. (2016)). These principles enable clinically faithful evaluation suitable for real-world deployment.

Evaluation Method. Each patient is associated with a sequence of T structured reports. The metric operates at the patient level and supports both single-report ($T = 1$) and sequential-report ($T > 1$) evaluations. In the **single-report** setting, evaluation is based on semantic and structural alignment, while in the **sequential-report** setting, temporal alignment is additionally incorporated to assess consistency across longitudinal disease trajectories. Formally, LUNGUESCORE evaluates similarity between predicted and gold reference sets of structured report findings as follows.

For each patient, we compare all predicted and gold reference findings across the entire sequence of reports. Let $\mathcal{S}^{\text{pred}} = (S_1^{\text{pred}}, \dots, S_T^{\text{pred}})$ and $\mathcal{S}^{\text{gold}} = (S_1^{\text{gold}}, \dots, S_T^{\text{gold}})$ denote the predicted and gold sequences for a given patient, where each $S_t^{(\cdot)}$ is the set of all structured findings at the t -th study. Pairwise similarity is computed over every possible pair of findings, pooled across all timepoints:

$$(f^{\text{pred}}, f^{\text{gold}}) \in \left(\bigcup_{t_p=1}^T S_{t_p}^{\text{pred}} \right) \times \left(\bigcup_{t_g=1}^T S_{t_g}^{\text{gold}} \right). \quad (1)$$

Each pair of findings is assigned a composite similarity score that captures alignment across semantic, temporal, and structural similarity dimensions, as defined below:

$$\text{MatchScore}(f^{\text{pred}}, f^{\text{gold}}) = \text{Semantic} \cdot (\text{Temporal if } T > 1) \cdot \text{Structural}. \quad (2)$$

Semantic similarity determines whether two findings express the same underlying clinical concept. Representation differs by setting: in single reports ($T = 1$), each finding is encoded as a linearized phrase combining the entity and all attributes (e.g., “opacity”-“left lung”-“nodular”-“slightly increased”); in sequential reports ($T > 1$), where findings must be tracked across time, we instead use ENTITYGROUP (Section 4). This enables lexically divergent but conceptually identical findings to align across multiple reports. Semantic similarity is then computed as the average cosine similarity between contextual embeddings from two domain-specific clinical BERT models—MedCPT and BioLORD (Jin et al. (2023); Remy et al. (2024))—selected for their ability to capture variability in chest X-ray language. Details of model selection are in Appendix C.3.

$$\text{Semantic}(f^{\text{pred}}, f^{\text{gold}}) = \text{cosine}(\text{Embed}(f^{\text{pred}}), \text{Embed}(f^{\text{gold}})) \quad (3)$$

324 **Temporal similarity** is defined only when $T > 1$ and captures alignment across timepoints. It
 325 ensures that findings are not only semantically similar but also temporally coherent with the patient’s
 326 disease progression. To prevent matches across unrelated timepoints, LUNGUESCORE prioritizes
 327 findings that occur in the same study timepoint t and TEMPORALGROUP. Temporal alignment
 328 receives the maximum score (= 1) when both study timepoint t and TEMPORALGROUP match, and a
 329 reduced score when only one matches, for example, when a predicted finding belongs to the correct
 330 TEMPORALGROUP but appears in a different study. Final scores are computed using equal weights:

$$\text{Temporal}(f^{\text{pred}}, f^{\text{gold}}) = w_S \cdot \mathbf{1}[S(f^{\text{pred}}) = S(f^{\text{gold}})] + w_G \cdot \mathbf{1}[G(f^{\text{pred}}) = G(f^{\text{gold}})]. \quad (4)$$

331 where S refers to the study timepoint t , G refers to the TEMPORALGROUP of findings across time,
 332 and equal weights ($w_S = w_G = 0.5$) are used in our implementation.

333 **Structural similarity** evaluates individual attributes (e.g. LOCATION, MEASUREMENT...) between
 334 predicted and gold reference findings, enabling fine-grained comparison. Each attribute is assigned a
 335 normalized weight $w_{\text{attribute}}$ based on its clinical importance, as determined by experts, reflecting its
 336 role in decision making (see Appendix C.1). Similarity is computed as:

$$\text{Structural}(f^{\text{pred}}, f^{\text{gold}}) = \sum_{\text{attribute}} w_{\text{attribute}} \cdot \text{sim}(f^{\text{pred}}[\text{attribute}], f^{\text{gold}}[\text{attribute}]), \quad (5)$$

341 where $\text{sim}(\cdot)$ returns 1 for exact matches on binary attributes³ and cosine similarity for non-binary
 342 attributes⁴ using the average of MedCPT and BioLORD contextual encoders. This ensures that
 343 evaluation captures both overall correctness and clinically critical attribute accuracy.

344 **Set-level matching with partial credit.** We can compute the combined MatchScore by multiplying
 345 semantic, temporal, and structural similarity scores (Equations 3-5), as shown in Equation 2. We then
 346 perform optimal bipartite matching between predicted findings i and gold reference findings j using
 347 MatchScore s_{ij} as edge weights, giving us sets of matched pairs $\{(f_m^{(\text{pred})}, f_n^{(\text{gold})})\}$, unmatched
 348 predicted findings $\{f_u^{(\text{pred})}\}$, and unmatched gold reference findings $\{f_v^{(\text{gold})}\}$. Matched pairs
 349 contribute similarity s_{mn} to true positives (TP), with residual $(1 - s_{mn})$ assigned to false positives
 350 (FP) and negatives (FN). Unmatched findings incur penalties based on their most similar finding:

$$\text{TP} = \sum_{(m,n)} s_{mn}, \text{FP} = \sum_{(m,n)} (1 - s_{mn}) + \sum_u \left(1 - \max_j s_{uj} \right), \text{FN} = \sum_{(m,n)} (1 - s_{mn}) + \sum_v \left(1 - \max_i s_{iv} \right). \quad (6)$$

355 This formulation supports **partial credit** based on alignment strength. Full credit is awarded only
 356 when a finding aligns simultaneously at the semantic, temporal, and structural levels. Partial matches
 357 contribute proportionally to the score, while unmatched findings in either set are penalized as FP or
 358 FN. This scoring scheme enables nuanced evaluation that distinguishes minor misalignments from
 359 complete misses. The final F1 score is computed from these TP, FP, and FN counts using the standard
 360 formula. Additional illustrative examples are provided in Appendix C.2.

363 6 EXPERIMENTS

364 We conduct three experiments from complementary perspectives: (1) performance of the structuring
 365 framework, (2) diagnostic utility of LUNGUESCORE as a single-report evaluation metric, and (3)
 366 benchmarking of single- and longitudinal-report generation models with LUNGUESCORE.

367 6.1 STRUCTURING FRAMEWORK VALIDATION

368 We evaluate the **structuring framework** on LUNGUESCORE, comprising 1,473 reports from 230 patients
 369 (1–15 studies each), including 30 patients with full longitudinal trajectories. Evaluation follows two
 370 stages: (i) single structuring, assessing localized semantic relations, and (ii) sequential structuring,
 371 evaluating consistency and organization of findings into clinical episodes across time.

372 ³Binary attributes: DXSTATUS (positive/negative) and DXCERTAINTY (definitive/tentative)

373 ⁴Non-binary attributes include: LOCATION, SEVERITY, ONSET, IMPROVED, WORSENED, PLACEMENT,
 374 NOCHANGE, MORPHOLOGY, DISTRIBUTION, MEASUREMENT, COMPARISON, PASTHX, OTHERSOURCE,
 375 ASSESSMENTLIMITATIONS

378 Table 1: Performance of various models under zero-shot and 5-shot settings. Left: single report
 379 structuring performance. Right: sequential report structuring performance.

381	382	383	Shot	Model	Single Structuring						Sequential Structuring					
					entity-relation			entity-relation-attribute			Entity Grouping			Temporal Grouping		
					F1	P	R	F1	P	R	F1	P	R	F1	P	R
384	385	386	Zero	GPT-4.1	0.91	0.83	1.00	0.78	0.79	0.77	0.67	0.68	0.71	0.84	0.83	0.86
				Qwen3	0.73	0.58	1.00	0.62	0.53	0.75	0.51	0.43	0.65	0.84	0.87	0.82
				Deepseek-v3	0.87	0.76	1.00	0.76	0.72	0.80	0.43	0.30	0.76	0.81	0.87	0.75
				Llama4-Maverick	0.81	0.68	1.00	0.69	0.64	0.76	0.37	0.24	0.77	0.60	0.87	0.47
				MedGemma-27b-text-it	0.75	0.59	1.00	0.20	0.28	0.16	0.33	0.30	0.37	0.74	0.85	0.66
				GPT-OSS-120b	0.92	0.85	1.00	0.70	0.71	0.69	0.62	0.57	0.69	0.83	0.86	0.80
389	390	391	5-shot	GPT-4.1	0.94	0.88	1.00	0.86	0.86	0.86	0.69	0.72	0.68	0.87	0.84	0.91
				Qwen3	0.92	0.85	1.00	0.84	0.83	0.85	0.64	0.58	0.71	0.85	0.87	0.84
				Deepseek-v3	0.93	0.88	1.00	0.86	0.85	0.86	0.67	0.61	0.75	0.86	0.89	0.84
				Llama4-Maverick	0.94	0.88	1.00	0.86	0.86	0.85	0.54	0.39	0.87	0.63	0.92	0.48
				MedGemma-27b-text-it	0.90	0.82	1.00	0.81	0.80	0.82	0.55	0.50	0.61	0.82	0.87	0.78
				GPT-OSS-120b	0.90	0.83	1.00	0.81	0.79	0.83	0.66	0.60	0.74	0.83	0.86	0.80

395
 396 **Single Structuring** We evaluate model performance on generating structured representations from
 397 individual reports by comparing predicted (*entity, relation, attribute*) triplets against expert annotations
 398 in LUNGUAGE. Using micro-averaged precision, recall, and F1 scores at both the entity–relation
 399 and full triplet levels, we assess GPT-4.1 Achiam et al. (2023) alongside several recent open-source
 400 LLMs Liu et al. (2024); OpenAI (2025); Sellergren et al. (2025); Touvron et al. (2023); Zheng
 401 et al. (2025) under the framework described in Section 4. As shown in Table 1, all models achieve
 402 perfect recall, with 5-shot prompting yielding F1 scores of 0.90–0.94 for entity–relation extraction and
 403 0.81–0.86 for full triplets. Performance improves further with more few-shot examples, demonstrating
 404 the robustness of the framework despite the schema’s complexity. Additional experiments, including
 405 vocabulary guidance, 10-shot prompting, and qualitative examples, are provided in Appendix B.3.

406
 407 **Sequential Structuring** The second stage evaluates whether models can group temporally dis-
 408 tributed findings into clinically meaningful categories, a task complicated by subtle semantic distinc-
 409 tions. For instance, “heart size” may group with “cardiomegaly,” whereas “mediastinal silhouette”
 410 concerns shape and can remain normal despite cardiomegaly. Using micro-averaged F1 scores,
 411 we observe that zero-shot prompting already yields strong temporal grouping performance (F1
 412 ≈ 0.80 –0.84 for GPT-4.1 and other LLMs), whereas entity grouping is noticeably more variable
 413 across models, particularly among open-source LLMs. Providing five in-context examples stabilizes
 414 the predictions and consistently improves entity grouping, with GPT-4.1 reaching an F1 of 0.69
 415 and most models exceeding 0.60, while temporal grouping remains high (F1 ≈ 0.82 –0.87). As
 416 detailed in Appendix B.4, the remaining discrepancies in entity grouping mainly concern how finely
 417 entities are grouped, for example whether closely related lexical variants or attribute-specific mentions
 418 are merged or split. Consequently, because LUNGUAGESCORE (Section 5, Equation 3) relies on
 419 continuous semantic, temporal, and structural similarity over matched findings rather than exact
 420 group identity, predictions that differ only in grouping granularity still receive high similarity when
 421 they follow the same diagnostic trajectory.

422 6.2 EVALUATING LUNGUAGESCORE ALIGNMENT WITH RADIOLOGIST JUDGMENTS

423
 424 We validate the diagnostic utility of LUNGUAGESCORE on the ReXVal dataset (Yu et al. (2023b)),
 425 which consists of 200 MIMIC-CXR report pairs annotated by six radiologists to benchmark alignment
 426 between automated metrics and expert judgments. As ReXVal contains only single reports, we
 427 evaluate the single-report version of LUNGUAGESCORE (semantic and structural alignment). We
 428 compare against BLEU, BERTScore, GREEN, FineRadScore, RaTEScore, RadGraph-F1, and
 429 RadGraph-XL F1, with implementation details in Appendix D.

430 Table 2 reports Kendall Tau and Pearson correlations between each metric and the number of
 431 radiologist-identified errors, where stronger alignment corresponds to more negative values. We also
 432 report 95% confidence intervals from 1,000 bootstrap resamples.

Our metric outperforms traditional structure- or semantics-based metrics (BLEU, BERTScore, RaTEScore, RadGraph-F1, RadGraph-XL F1) but falls slightly short of LLM-derived scores (FineRadScore, GREEN), which are explicitly tuned to ReXVal’s error taxonomy. Nonetheless, LUNGScore achieves performance close to these metrics while relying only on semantic and structural alignment rather than error-type supervision. Additional analyses in Appendix D show that LUNGScore correlates strongly with all other metrics.

6.3 BENCHMARKING SINGLE-REPORT AND SEQUENTIAL REPORT GENERATION MODELS

We further validate LUNGScore by benchmarking it against existing evaluation methods across a diverse set of report generation models, assessing its capacity to capture clinically meaningful differences at both single-report and patient-level scales. We categorize the models based on input modality: those utilizing only the *current timewindow single image* (Cvt2DistilGPT2 (Nicolson et al., 2023), RGRG (Tanida et al., 2023), MedGemma (Sellergren et al., 2025), Lingshu (Xu et al., 2025), CheXAgent (Chen et al., 2024)), and those incorporating the *current image plus prior context* (e.g., history section or prior image) (Medversa (Zhou et al., 2024), LIBRA (Zhang et al., 2025), MAIRA-2 (Bannur et al., 2024)).

Radiology report generation All models require frontal chest X-rays. MAIRA-2 additionally uses lateral images when available. RGRG and Cvt2DistilGPT2 generate findings sections, while Medversa, MedGemma, Lingshu, CheXAgent, and LIBRA produce full reports that include both findings and impression; we standardize all outputs into complete reports. The *single+prior* group is configured to consume both the current and a prior context. MAIRA-2 (*standard*) and Medversa also incorporate the history/indication text as contextual input. Implementation details are provided in Appendix E.

Single-report setting In this setting, we compare generated reports with ground-truth references on a study-by-study basis using the same patient as in the sequential evaluation, restricted to the 67 studies that contain frontal images.⁵ Reference reports combine findings and impression. Table 3 summarizes results across metrics, including LUNGScore. For LUNGScore, we use LUNGScore as ground truth and compare against outputs from the structuring framework in Section 4. Overall, models that use both current and prior context perform better. MAIRA-2 (standard, highest-performing), Medversa, and LIBRA achieve higher scores than the single-image baselines on LUNGScore (single), indicating a clear benefit from incorporating longitudinal context even when evaluating single reports. Notably, while their advantage over single-image models is relatively modest under existing metrics, LUNGScore reveals a clearer gap between context-aware and single-image systems.

Sequential Setting We use the same reports as in the single-report setting but additionally include the history/indication section to provide context for patient trajectories. All models are evaluated in this sequential setting, including those that only take the current image as input. This is because the benchmark is organized around patient-level sequences: in routine practice, radiologists interpret each study in light of previous examinations, and the resulting reports describe how findings evolve over time rather than isolated snapshots. As a consequence, the reference reports form a temporally coherent longitudinal narrative for each patient. Running any model independently at each timewindow therefore induces its own predicted trajectory, which LUNGScore can meaningfully compare against this coherent reference sequence.

Table 2: Kendall Tau and Pearson correlation coefficients (with 95% CIs) between single-report metrics and the total number of radiologist-annotated errors in each report, across the ReXVal dataset. Note that FineRadScore was inverted for comparability.

Metric	Kendall Tau	Pearson
BLEU	-0.38 (-0.29, -0.48)	-0.53 (-0.45, -0.62)
BERTScore	-0.50 (-0.43, -0.57)	-0.63 (-0.55, -0.70)
GREEN	-0.63 (-0.56, -0.69)	-0.73 (-0.67, -0.78)
1/FineRadScore	-0.69 (-0.63, -0.74)	-0.75 (-0.69, -0.80)
RaTEScore	-0.52 (-0.44, -0.59)	-0.63 (-0.55, -0.70)
RadGraph F1	-0.57 (-0.50, -0.63)	-0.68 (-0.62, -0.74)
RadGraph-XL F1	-0.53 (-0.45, -0.62)	-0.63 (-0.56, -0.72)
LUNGScore	-0.58 (-0.52, -0.64)	-0.69 (-0.63, -0.74)

⁵Studies without frontal images were excluded, which created gaps in sequential analyses for 5 of 10 patients.

486
487
488
489 Table 3: Evaluation of radiology report generation models using multiple metrics. Scores are averages
490 with 95% CIs.

490 Input	491 Model	492 Single-report setting					493 Sequential
		494 RaTEScore	495 GREEN	496 1/FineRadScore	497 RadGraph F1	498 LUNGUESCORE	
499 single	Cvt2DistilGPT2	0.491 (0.46, 0.52)	0.240 (0.19, 0.29)	0.167 (0.14, 0.20)	0.179 (0.15, 0.21)	0.367 (0.34, 0.40)	0.371 (0.33, 0.41)
	RGRG	0.547 (0.53, 0.57)	0.266 (0.23, 0.30)	0.139 (0.11, 0.17)	0.264 (0.23, 0.29)	0.406 (0.38, 0.43)	0.391 (0.36, 0.42)
	MedGemma	0.495 (0.48, 0.51)	0.149 (0.10, 0.20)	0.127 (0.11, 0.14)	0.133 (0.12, 0.15)	0.318 (0.30, 0.34)	0.345 (0.32, 0.37)
	Lingshu	0.483 (0.46, 0.50)	0.173 (0.13, 0.22)	0.141 (0.11, 0.17)	0.150 (0.13, 0.18)	0.344 (0.32, 0.37)	0.356 (0.33, 0.38)
	CheXAgent	0.528 (0.50, 0.55)	0.241 (0.20, 0.29)	0.131 (0.12, 0.14)	0.228 (0.20, 0.26)	0.380 (0.35, 0.41)	0.388 (0.36, 0.42)
500 single + prior	Medversa	0.543 (0.52, 0.57)	0.314 (0.26, 0.37)	0.183 (0.15, 0.22)	0.238 (0.21, 0.27)	0.409 (0.38, 0.44)	0.410 (0.37, 0.45)
	LIBRA	0.526 (0.50, 0.55)	0.266 (0.22, 0.30)	0.127 (0.12, 0.14)	0.227 (0.20, 0.26)	0.414 (0.38, 0.45)	0.417 (0.38, 0.43)
	MAIRA-2 (standard)	0.564 (0.54, 0.59)	0.325 (0.28, 0.37)	0.156 (0.14, 0.18)	0.274 (0.25, 0.30)	0.429 (0.40, 0.46)	0.432 (0.41, 0.46)
	MAIRA-2 (cascade)	0.547 (0.53, 0.57)	0.299 (0.25, 0.34)	0.171 (0.13, 0.21)	0.233 (0.21, 0.26)	0.419 (0.39, 0.45)	0.416 (0.38, 0.45)

501 As shown in Table 3, models that leverage prior context (*single + prior*) rank above the single-image
502 models, with MAIRA-2 achieving the best performance and Medversa second. Models that omit
503 this prior context (Cvt2DistilGPT2, RGRG, MedGemma, Lingshu, CheXAgent) perform worse;
504 notably, when comparing the single-report and sequential evaluations, the LUNGUESCORE scores
505 of RGRG and MAIRA-2 (cascade) decrease, whereas the scores of the other models improve. This
506 pattern indicates that models which appear strong in the single-report setting can still produce
507 temporally inconsistent diagnostic trajectories once their predictions are examined across the full
508 patient sequence. LUNGUESCORE makes these temporal inconsistencies explicit and highlights
509 the importance of evaluating report generators in the sequential setting. Further analysis of error
510 sensitivity is provided in Section D.

511 7 CONCLUSION

512 This work introduces a comprehensive pipeline for evaluating radiology reports, grounded in LUN-
513 GUAGE, a fine-grained benchmark for single and sequential structured chest X-ray reports. To our
514 knowledge, it is the first benchmark and evaluation framework explicitly designed for longitudinal
515 chest X-ray report generation and patient-level structured assessment. The dataset is intentionally
516 designed as a dense, expert-verified evaluation resource, comprising 1,473 single reports and longi-
517 tudinal patient trajectories with rich entity–attribute structure and temporal alignments curated by
518 board-certified radiologists. Building on this foundation, we propose a two-stage LLM-based struc-
519 turing framework that reliably maps free-text reports into schema-aligned representations across both
520 single and sequential settings, and LUNGUESCORE, a clinically grounded metric that evaluates
521 model outputs along semantic, structural, and temporal dimensions.

522 LUNGUESCORE operates on entity-centered representations that bundle attributes and temporal
523 links for each finding, enabling joint assessment of diagnostic status, spatial and descriptive attributes,
524 longitudinal change, and relevant context. Reports are first mapped by an LLM into schema-aligned
525 structures, then evaluated by a fixed scoring function over semantic, temporal, and attribute fields,
526 yielding transparent, attribute-wise interpretable scores that remain stable under small errors in
527 the structured inputs. Empirically, LUNGUESCORE correlates well with radiologist-annotated
528 errors on ReXVal in the single-report setting and, in the longitudinal setting, more clearly separates
529 context-aware models from single-image baselines while reducing penalties for clinically appropriate
530 but textually omitted findings. We expect our work to serve as a practical, interpretable testbed for
531 fine-grained single-report and longitudinal evaluation, and we discuss remaining limitations and
532 avenues for extension in Appendix F.

533
534
535
536
537
538
539

540 REFERENCES
541

542 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman,
543 Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report.
544 *arXiv preprint arXiv:2303.08774*, 2023.

545 Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jin, Tristan Naumann, and
546 Matthew McDermott. Publicly available clinical BERT embeddings. In *Proceedings of the 2nd*
547 *Clinical Natural Language Processing Workshop*, pp. 72–78, Minneapolis, Minnesota, USA,
548 June 2019. Association for Computational Linguistics. doi: 10.18653/v1/W19-1909. URL
549 <https://www.aclweb.org/anthology/W19-1909>.

550 Satanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved
551 correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic*
552 *evaluation measures for machine translation and/or summarization*, pp. 65–72, 2005.

553 Alexander A Bankier, Heber MacMahon, Thomas Colby, Pierre Alain Gevenois, Jin Mo Goo, Ann NC
554 Leung, David A Lynch, Cornelia M Schaefer-Prokop, Noriyuki Tomiyama, William D Travis, et al.
555 Fleischner society: glossary of terms for thoracic imaging. *Radiology*, 310(2):e232558, 2024.

556 Shruthi Bannur, Kenza Bouzid, Daniel C Castro, Anton Schwaighofer, Anja Thieme, Sam Bond-
557 Taylor, Maximilian Ilse, Fernando Pérez-García, Valentina Salvatelli, Harshita Sharma, et al.
558 MAIRA-2: Grounded radiology report generation. *arXiv preprint arXiv:2406.04449*, 2024.

559 Olivier Bodenreider. The unified medical language system (umls): integrating biomedical terminology.
560 *Nucleic acids research*, 32(suppl_1):D267–D270, 2004.

561 Felix Busch, Lena Hoffmann, Daniel Pinto Dos Santos, Marcus R Makowski, Luca Saba, Philipp
562 Prucker, Martin Hadamitzky, Nassir Navab, Jakob Nikolas Kather, Daniel Truhn, et al. Large
563 language models for structured reporting in radiology: past, present, and future. *European*
564 *Radiology*, pp. 1–14, 2024.

565 Wendy W Chapman, Prakash M Nadkarni, Lynette Hirschman, Leonard W D’avolio, Guergana K
566 Savova, and Ozlem Uzuner. Overcoming barriers to nlp for clinical text: the role of shared tasks
567 and the need for additional creative solutions, 2011.

568 Zhihong Chen, Maya Varma, Justin Xu, Magdalini Paschali, Dave Van Veen, Andrew Johnston, Alaa
569 Youssef, Louis Blankemeier, Christian Blauthgen, Stephan Altmayer, et al. A vision-language
570 foundation model to enhance efficiency of chest x-ray interpretation. *arXiv e-prints*, pp. arXiv-
571 2401, 2024.

572 Jean-Benoit Delbrouck, Pierre Chambon, Christian Blauthgen, Emily Tsai, Omar Almusa, and Curtis
573 Langlotz. Improving the factual correctness of radiology report generation with semantic rewards.
574 In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pp. 4348–4360, Abu
575 Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. URL
576 <https://aclanthology.org/2022.findings-emnlp.319>.

577 Jean-Benoit Delbrouck, Pierre Chambon, Zhihong Chen, Maya Varma, Andrew Johnston, Louis
578 Blankemeier, Dave Van Veen, Tan Bui, Steven Truong, and Curtis Langlotz. RadGraph-XL:
579 A large-scale expert-annotated dataset for entity and relation extraction from radiology reports.
580 In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association for*
581 *Computational Linguistics ACL 2024*, pp. 12902–12915, Bangkok, Thailand and virtual meeting,
582 August 2024. Association for Computational Linguistics. URL <https://aclanthology.org/2024.findings-acl.765>.

583 Dina Demner-Fushman, Wendy W Chapman, and Clement J McDonald. What can natural language
584 processing do for clinical decision support? *Journal of biomedical informatics*, 42(5):760–772,
585 2009.

586 Felix J Dorfner, Liv Jürgensen, Leonhard Donle, Fares Al Mohamad, Tobias R Bodenmann, Mason C
587 Cleveland, Felix Busch, Lisa C Adams, James Sato, Thomas Schultz, et al. Comparing commercial
588 and open-source large language models for labeling chest radiograph reports. *Radiology*, 313(1):
589 e241139, 2024.

594 Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann,
 595 Jianfeng Gao, and Hoifung Poon. Domain-specific language model pretraining for biomedical
 596 natural language processing, 2020.

597 Iryna Hartsock, Cyrillo Araujo, Les Folio, and Ghulam Rasool. Improving radiology report conciseness
 598 and structure via local large language models. *Journal of Imaging Informatics in Medicine*,
 599 pp. 1–12, 2025.

600 Alyssa Huang, Oishi Banerjee, Kay Wu, Eduardo Pontes Reis, and Pranav Rajpurkar. Fineradscore:
 601 A radiology report line-by-line evaluation technique generating corrections with severity scores.
 602 *arXiv preprint arXiv:2405.20613*, 2024.

603 Saahil Jain, Ashwin Agrawal, Adriel Saporta, Steven QH Truong, Du Nguyen Duong, Tan Bui, Pierre
 604 Chambon, Yuhao Zhang, Matthew P Lungren, Andrew Y Ng, et al. Radgraph: Extracting clinical
 605 entities and relations from radiology reports. *arXiv preprint arXiv:2106.14463*, 2021.

606 Qiao Jin, Won Kim, Qingyu Chen, Donald C Comeau, Lana Yeganova, W John Wilbur, and Zhiyong
 607 Lu. Medcpt: Contrastive pre-trained transformers with large-scale pubmed search logs for zero-shot
 608 biomedical information retrieval. *Bioinformatics*, 39(11):btad651, 2023.

609 Alistair EW Johnson, Tom J Pollard, Seth J Berkowitz, Nathaniel R Greenbaum, Matthew P Lungren,
 610 Chih-ying Deng, Roger G Mark, and Steven Horng. Mimic-cxr, a de-identified publicly available
 611 database of chest radiographs with free-text reports. *Scientific data*, 6(1):317, 2019.

612 Sameer Khanna, Adam Dejl, Kibo Yoon, Steven QH Truong, Hanh Duong, Agustina Saenz, and
 613 Pranav Rajpurkar. Radgraph2: Modeling disease progression in radiology reports via hierarchical
 614 information extraction. In *Machine Learning for Healthcare Conference*, pp. 381–402. PMLR,
 2023.

615 Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo
 616 Kang. Biobert: a pre-trained biomedical language representation model for biomedical text mining.
 617 *Bioinformatics*, 36(4):1234–1240, 2020.

618 Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization
 branches out*, pp. 74–81, 2004.

619 Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao,
 620 Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint
 621 arXiv:2412.19437*, 2024.

622 Xiaohong Liu, Hao Liu, Guoxing Yang, Zeyu Jiang, Shuguang Cui, Zhaoze Zhang, Huan Wang,
 623 Liyuan Tao, Yongchang Sun, Zhu Song, et al. A generalist medical language model for disease
 624 diagnosis assistance. *Nature Medicine*, pp. 1–11, 2025.

625 Stéphane M Meystre, Guergana K Savova, Karin C Kipper-Schuler, and John F Hurdle. Extracting
 626 information from textual documents in the electronic health record: a review of recent research.
 627 *Yearbook of medical informatics*, 17(01):128–144, 2008.

628 Aaron Nicolson, Jason Dowling, and Bevan Koopman. Improving chest x-ray report generation by
 629 leveraging warm starting. *Artificial intelligence in medicine*, 144:102633, 2023.

630 OpenAI. gpt-oss-120b gpt-oss-20b model card, 2025. URL <https://arxiv.org/abs/2508.10925>.

631 Sophie Ostmeier, Justin Xu, Zhihong Chen, Maya Varma, Louis Blankemeier, Christian Bluethgen,
 632 Arne Edward Michalson, Michael Moseley, Curtis Langlotz, Akshay S Chaudhari, et al. Green:
 633 Generative radiology report evaluation and error notation. *arXiv preprint arXiv:2405.03595*, 2024.

634 Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic
 635 evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association
 636 for Computational Linguistics*, pp. 311–318, 2002.

637 Ewoud Pons, Loes MM Braun, MG Myriam Hunink, and Jan A Kors. Natural language processing
 638 in radiology: a systematic review. *Radiology*, 279(2):329–343, 2016.

648 Vishwanatha M Rao, Serena Zhang, Julian N Acosta, Subathra Adithan, and Pranav Rajpurkar.
 649 Rexerr: Synthesizing clinically meaningful errors in diagnostic radiology reports. In *Biocomputing*
 650 *2025: Proceedings of the Pacific Symposium*, pp. 70–81. World Scientific, 2024.

651 Vishwanatha M Rao, Serena Zhang, Julian N Acosta, Subathra Adithan, and Pranav Rajpurkar.
 652 Rexerr-v1: Clinically meaningful chest x-ray report errors derived from mimic-cxr (version 1.0.0).
 653 *Physionet*, 2025. doi: <https://doi.org/10.13026/9dns-vd94>.

654 François Remy, Kris Demuynck, and Thomas Demeester. Biolord-2023: semantic textual representations
 655 fusing large language models and clinical knowledge graph insights. *Journal of the American
 656 Medical Informatics Association*, 31(9):1844–1855, 2024.

657 Guergana K Savova, James J Masanz, Philip V Ogren, Jiaping Zheng, Sunghwan Sohn, Karin C
 658 Kipper-Schuler, and Christopher G Chute. Mayo clinical text analysis and knowledge extraction
 659 system (ctakes): architecture, component evaluation and applications. *Journal of the American
 660 Medical Informatics Association*, 17(5):507–513, 2010.

661 Andrew Sellergren, Sahar Kazemzadeh, Tiam Jaroensri, Atilla Kiraly, Madeleine Traverse, Timo
 662 Kohlberger, Shawn Xu, Fayaz Jamil, Cian Hughes, Charles Lau, et al. Medgemma technical report.
 663 *arXiv preprint arXiv:2507.05201*, 2025.

664 Akshay Smit, Saahil Jain, Pranav Rajpurkar, Anuj Pareek, Andrew Y Ng, and Matthew P Lungren.
 665 Chexbert: combining automatic labelers and expert annotations for accurate radiology report
 666 labeling using bert. *arXiv preprint arXiv:2004.09167*, 2020.

667 Tim Tanida, Philip Müller, Georgios Kaassis, and Daniel Rueckert. Interactive and explainable region-
 668 guided radiology report generation. In *Proceedings of the IEEE/CVF Conference on Computer
 669 Vision and Pattern Recognition*, pp. 7433–7442, 2023.

670 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
 671 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and
 672 efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.

673 Yanshan Wang, Liwei Wang, Majid Rastegar-Mojarad, Sungrim Moon, Feichen Shen, Naveed Afzal,
 674 Sijia Liu, Yuqun Zeng, Saeed Mehrabi, Sunghwan Sohn, et al. Clinical information extraction
 675 applications: a literature review. *Journal of biomedical informatics*, 77:34–49, 2018.

676 Piotr Woźnicki, Caroline Laqua, Ina Fiku, Amar Hekalo, Daniel Truhn, Sandy Engelhardt, Jakob
 677 Kather, Sebastian Foersch, Tugba Akinci D’Antonoli, Daniel Pinto dos Santos, et al. Automatic
 678 structuring of radiology reports with on-premise open-source large language models. *European
 679 Radiology*, pp. 1–12, 2024.

680 Joy T Wu, Nkechinyere N Agu, Ismini Lourentzou, Arjun Sharma, Joseph A Paguio, Jasper S
 681 Yao, Edward C Dee, William Mitchell, Satyananda Kashyap, Andrea Giovannini, et al. Chest
 682 imangenome dataset for clinical reasoning. *arXiv preprint arXiv:2108.00316*, 2021.

683 Weiwen Xu, Hou Pong Chan, Long Li, Mahani Aljunied, Ruifeng Yuan, Jianyu Wang, Chenghao
 684 Xiao, Guizhen Chen, Chaoqun Liu, Zhaodonghui Li, et al. Lingshu: A generalist foundation model
 685 for unified multimodal medical understanding and reasoning. *arXiv preprint arXiv:2506.07044*,
 686 2025.

687 Feiyang Yu, Mark Endo, Rayan Krishnan, Ian Pan, Andy Tsai, Eduardo Pontes Reis, Eduardo Kaiser
 688 Ururahy Nunes Fonseca, Henrique Min Ho Lee, Zahra Shakeri Hossein Abad, Andrew Y Ng, et al.
 689 Evaluating progress in automatic chest x-ray radiology report generation. *Patterns*, 4(9), 2023a.

690 Feiyang Yu, Mark Endo, Rayan Krishnan, Ian Pan, Andy Tsai, Eduardo Pontes Reis, EKU Fonseca,
 691 Henrique Lee, Zahra Shakeri, Andrew Ng, et al. Radiology report expert evaluation (rexval)
 692 dataset, 2023b.

693 Juan Manuel Zambrano Chaves, Shih-Cheng Huang, Yanbo Xu, Hanwen Xu, Naoto Usuyama,
 694 Sheng Zhang, Fei Wang, Yujia Xie, Mahmoud Khademi, Ziyi Yang, et al. A clinically accessible
 695 small multimodal radiology model and evaluation metric for chest x-ray findings. *Nature
 696 Communications*, 16(1):3108, 2025.

702 Mengliang Zhang, Xinyue Hu, Lin Gu, Tatsuya Harada, Kazuma Kobayashi, Ronald Summers, and
703 Yingying Zhu. Cad-chest: Comprehensive annotation of diseases based on mimic-cxr radiology
704 report. (*No Title*), 2023.

705 Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. Bertscore: Evaluating
706 text generation with bert. *arXiv preprint arXiv:1904.09675*, 2019.

708 Xi Zhang, Zaiqiao Meng, Jake Lever, and Edmond SL Ho. Libra: Leveraging temporal images for
709 biomedical radiology analysis. In *Findings of the Association for Computational Linguistics: ACL*
710 2025, pp. 17275–17303, 2025.

711 Weike Zhao, Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. Ratescore: A
712 metric for radiology report generation. *arXiv preprint arXiv:2406.16845*, 2024.

714 Xingyu Zheng, Yuye Li, Haoran Chu, Yue Feng, Xudong Ma, Jie Luo, Jinyang Guo, Haotong Qin,
715 Michele Magno, and Xianglong Liu. An empirical study of qwen3 quantization. *arXiv preprint*
716 *arXiv:2505.02214*, 2025.

717 Hong-Yu Zhou, Subathra Adithan, Julián Nicolás Acosta, Eric J Topol, and Pranav Rajpurkar. A
718 generalist learner for multifaceted medical image interpretation. *arXiv preprint arXiv:2405.07988*,
719 2024.

721

722

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756	SUPPLEMENTARY CONTENTS	
757		
758		
759	A LUNGUAGE Details	16
760	A.1 Single-report Schema: Entity and Relation Definition	17
761	A.2 Task-specific Vocabulary Construction	19
762	A.3 Annotation Process	25
763	A.4 Detailed Comparison with Existing Structured-Report Datasets	32
764		
765		
766	B Two-stage Structuring Framework details	34
767	B.1 Single Structuring Prompt	34
768	B.2 Sequential Structuring Prompt	36
769	B.3 Single Structuring Analysis	37
770	B.4 Sequential Structuring Analysis	38
771		
772		
773		
774	C LUNGUAGESCORE Details	41
775	C.1 Attribute Weights of LUNGUAGESCORE	41
776	C.2 LUNGUAGESCORE examples	41
777	C.3 Clinical BERT Model Selection	44
778		
779		
780		
781	D Metric Validation	45
782		
783	E Synthetic Report Generation Details	49
784		
785	F Limitations and Future Directions	50
786		
787		
788		
789		
790		
791		
792		
793		
794		
795		
796		
797		
798		
799		
800		
801		
802		
803		
804		
805		
806		
807		
808		
809		

810 A LANGUAGE DETAILS

812 **Dataset preparation** LUNGUAGE aims to support patient-level evaluation of chest X-ray reports
 813 by modeling longitudinal diagnostic scenarios. To this end, we curated a benchmark dataset from the
 814 official test split of MIMIC-CXR, selecting patients with between 1 and 15 sequential studies. This
 815 yielded 230 patients with a total of 1,473 reports.

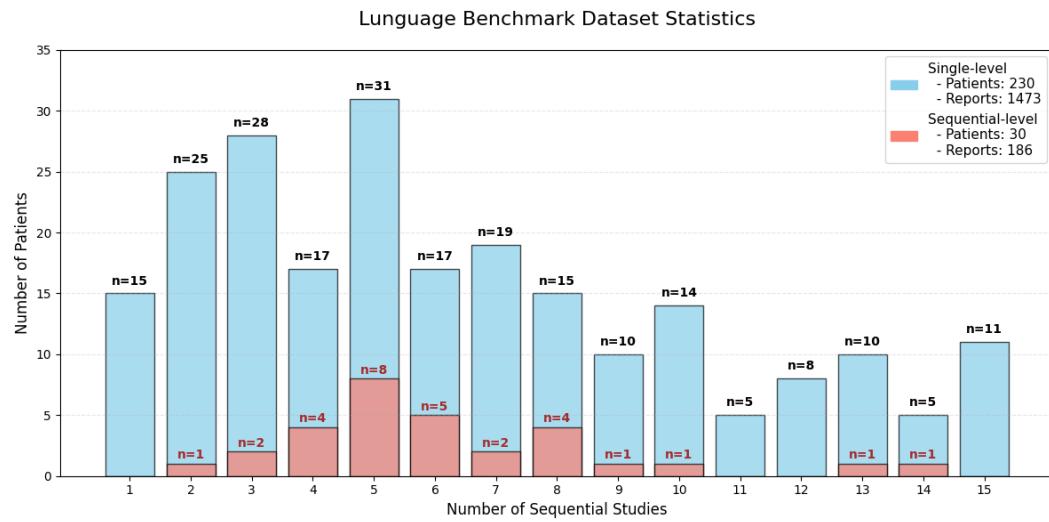
816 We followed the official MIMIC-CXR preprocessing protocol to extract structured text from each
 817 report. Specifically, we parsed the history (including “Indication”), findings, and impression
 818 sections. The history/indication field provides contextual information relevant to diagnostic reasoning,
 819 such as presenting symptoms (e.g., “fever,” “fatigue,” “cough”) or evaluation intents (e.g., “rule out
 820 pneumonia”). In contrast, the findings and impression sections describe image-based observations
 821 and interpretations.

822 Section-level coverage across the dataset is summarized as:

- 824 • **History (i.e., Indication):** 1,362 reports (92.5%)
- 825 • **Findings:** 1,224 reports (83.1%)
- 826 • **Impression:** 1,015 reports (68.9%)

828 Among the reports, 767 contained both findings and impression sections, 457 had findings only, 248
 829 had impression only, and 1 contained only a history section. We excluded infrequently occurring
 830 sections such as comparison (often containing anonymized metadata using placeholders like “__”),
 831 and technique (e.g., “AP view”), as these appeared in fewer than 5% of cases and were not directly
 832 relevant to diagnostic content.

833 To preserve diagnostic integrity and linguistic variability, we retained all reports in their original
 834 form without content filtering. This includes templated reports (e.g., “No acute cardiopulmonary
 835 process”) and incomplete notes. All reports were annotated using our schema-based pipeline with no
 836 preprocessing beyond section parsing. Structured reports were constructed by directly using the raw
 837 textual expressions from the original reports, rather than replacing them with normalized terms, to
 838 maintain alignment with the radiologists’ source language.



858 **Figure A.1: Distribution of the number of imaging studies per patient in LUNGUAGE.** Skyblue
 859 bars indicate the number of patients for each trajectory length (i.e., number of chest X-ray studies),
 860 reflecting the single-report annotation coverage. Salmon bars represent the subset of patients whose
 861 reports are also annotated at the longitudinal level. Values above the bars show the number of patients
 862 per group (n =), and for salmon bars, the number of patients with sequential annotations. The legend
 863 summarizes the total number of patients and reports included at each annotation level.

864 A.1 SINGLE-REPORT SCHEMA: ENTITY AND RELATION DEFINITION
865866 LANGUAGE represents each radiology report as a structured collection of (*entity*, *relation*, *attribute*)
867 triplets. This schema is designed to encode the diagnostic content of reports in a form that supports
868 structured analysis, longitudinal reasoning, and machine-readable interpretation. It captures both
869 observable features from chest X-ray (CXR) images and additional contextual elements embedded in
870 clinical narratives.871 A.1.1 ENTITY TYPES
872873 Entities represent clinically meaningful units such as findings, diagnoses, objects, or background
874 context. Each entity is assigned one of six mutually exclusive Cat (category) labels, depending on
875 whether it originates from the CXR image or external clinical sources. These six labels fall into two
876 broad groups:

- 878 • **Chest X-ray Findings** are entities that can be directly visualized on the chest X-ray or
879 inferred through image-based interpretation, possibly with minimal supporting context.
880 These form the core of radiologic description and are divided into the following types:
 - 881 – **PF (Perceptual Findings)**: Visual features that are explicitly visible in the image and
882 correspond to anatomical or pathological structures (e.g., “opacity”, “pleural effusion”,
883 “pneumothorax”). These are the most direct and objective form of image evidence.
 - 884 – **CF (Contextual Findings)**: Diagnoses that require interpretation of visual findings in
885 light of limited contextual knowledge (e.g., “pneumonia”, “congestive heart failure”).
886 These may involve reasoning beyond the image but still rely primarily on radiographic
887 evidence.
 - 888 – **OTH (Other Objects)**: Non-anatomic elements such as medical devices, surgical
889 hardware, or foreign materials visible on the image (e.g., “endotracheal tube”, “central
890 venous catheter”, “foreign body”). These often require placement verification or
891 complication monitoring.
- 892 • **Non Chest X-ray Findings** are entities that cannot be determined from the image alone and
893 must be inferred from patient history, clinical documentation, or other diagnostic modalities:
 - 894 – **COF (Clinical Objective Findings)**: Structured clinical measurements or physical
895 findings derived from sources such as laboratory tests or vital signs (e.g., “elevated
896 white cell count”, “low oxygen saturation”). These provide objective support for
897 contextual interpretation.
 - 898 – **NCD (Non-CXR Diagnosis)**: Diagnoses that originate from non-CXR modalities
899 (e.g., CT, MRI, serology) and are either mentioned for completeness or used to explain
900 findings (e.g., “stroke”, “AIDS”).
 - 901 – **PATIENT INFO**: Historical or subjective patient information, such as symptoms
902 or clinical background, that contributes to interpretation (e.g., “fever”, “history of
903 malignancy”, “recent trauma”).

904 Each entity is additionally annotated with the following attributes that define its diagnostic interpreta-
905 tion within the report:

- 906 • **DxStatus**: Indicates whether the entity is considered present or absent in the current study.
907 This label is determined from report language and includes implications from stability or
908 change. For example, “resolved effusion” is annotated as Positive, while “unchanged
909 opacity” is Positive unless the prior state was normal, in which case it is Negative.
- 910 • **DxCertainty**: Reflects the level of confidence expressed by the radiologist, labeled as either
911 Definitive or Tentative. Typical cues include phrases like “suggests”, “cannot exclude”,
912 or “possibly indicative of”, all leading to a tentative label.

914 A.1.2 RELATION TYPES
915916 Relations describe either attributes of a single entity or clinically relevant links between multiple
917 entities. All relations must be grounded in the report text and can span across sentences within the
918 same section.

918 **1. Diagnostic Reasoning** These relations connect semantically and clinically related entities. They
 919 encode the logic behind diagnostic interpretation.
 920

- 921 • **Associate:** A bidirectional, non-causal relationship between entities that co-occur or are
 922 conceptually linked (e.g., “opacity” ↔ “consolidation”). When Evidence is used, a corre-
 923 sponding Associate is also required in the reverse direction.
- 924 • **Evidence:** A unidirectional relation in which a finding supports a diagnosis (e.g., “pneumo-
 925 nia” → “opacity”).

927 **2. Spatial and Descriptive Attributes** These relations describe intrinsic visual characteristics of an
 928 entity as observed within a single chest X-ray image. Unlike temporal attributes, these do not require
 929 comparison with prior studies. Instead, they provide descriptive detail that refines the interpretation
 930 of a finding or object in terms of location, form, extent, intensity, and symmetry.

- 931 • **Location:** Specifies the anatomical or spatial position of the entity (e.g., “right upper lobe”,
 932 “carina above 3 cm”). An entity may have multiple location labels, annotated as a comma-
 933 separated list (e.g., “right upper lobe, suprahilar”). Location applies to both disease findings
 934 and device placements (e.g., “fragmentation” of “sternal wires”).
- 935 • **Morphology:** Describes the shape, form, or structural appearance of the entity (e.g., “nodu-
 936 lar”, “linear”, “reticular”, “confluent”). Morphological terms help differentiate types of
 937 opacities or identify characteristic patterns of pathology.
- 938 • **Distribution:** Refers to the anatomical spread or pattern of the entity (e.g., “focal”, “dif-
 939 fuse”, “multifocal”, “bilateral”). This helps characterize whether the finding is localized or
 940 widespread, and whether it follows typical anatomical distributions.
- 941 • **Measurement:** Captures quantitative properties such as size, count, or volume (e.g., “2.5
 942 cm”, “few”, “multiple”). These descriptors are typically numerical or ordinal and assist in
 943 severity grading or follow-up comparison.
- 944 • **Severity:** Reflects the degree of abnormality or clinical impact, often based on radiologic
 945 intensity or extent (e.g., “mild”, “moderate”, “severe”, “marked”).
- 946 • **Comparison:** Indicates asymmetry or difference across anatomical sides or regions within
 947 the same image (e.g., “left greater than right”, “right lung appears denser”). This is distinct
 948 from temporal comparison and only refers to spatial contrasts visible in the current image.

950 **3. Temporal Change** These relations capture how an entity has changed over time by comparing the
 951 current study to previous imaging or known clinical baselines. Temporal attributes are essential for
 952 longitudinal interpretation and reflect disease progression, treatment response, or clinical stability.
 953 Unlike static descriptors, these attributes require temporal context and often imply clinical decision
 954 points.

- 955 • **Onset:** Indicates the timing or duration of a finding as described in the report (e.g., “acute”,
 956 “subacute”, “chronic”, “new”). These descriptors suggest whether a condition has recently
 957 appeared or has been long-standing.
- 958 • **Improved:** Signals that a finding has regressed or resolved compared to a prior state (e.g.,
 959 “resolved effusion”, “decreased consolidation”). It is typically associated with positive
 960 treatment response or natural recovery.
- 961 • **Worsened:** Indicates that the condition has progressed, increased in extent, or become more
 962 severe over time (e.g., “enlarging opacity”, “increased pleural effusion”). This is often
 963 associated with disease progression or complications.
- 964 • **No Change:** Describes a finding that has remained stable since a prior study (e.g., “un-
 965 changed opacity”, “persistent nodule”). Although these are annotated as Positive by
 966 default, they are marked as Negative if the prior state was normal (i.e., continued absence
 967 of disease).
- 968 • **Placement:** Applies specifically to entities labeled as OTH (devices). It describes both the
 969 position (e.g., “in expected position”, “malpositioned”) and temporal actions involving the
 970 device (e.g., “inserted”, “withdrawn”, “removed”). This attribute is crucial for monitoring
 971 device-related interventions over time.

972 **4. Contextual Information** This category captures auxiliary information that influences the interpretation
 973 of findings but is not a primary descriptor of the radiologic appearance. These relations provide
 974 critical contextual cues—such as modality constraints, patient factors, or historical references—that
 975 support diagnostic interpretation. While not visual in the conventional sense, they are essential for
 976 accurately situating radiologic findings within the broader clinical scenario.

977 • **Past Hx:** Refers to the patient’s prior medical or surgical history that contextualizes current
 978 findings (e.g., “status post lobectomy”, “known tuberculosis”). These mentions often justify
 979 or explain current observations or exclude certain diagnoses.

980 • **Other Source:** Indicates that part of the reported information is derived from modalities
 981 other than chest X-ray (e.g., “seen on CT”, “confirmed on MRI”). This distinction is
 982 important when findings cannot be visualized directly on the image being interpreted.

983 • **Assessment Limitations:** Describes technical or procedural factors that constrain the
 984 radiologist’s ability to interpret the image accurately (e.g., “poor inspiration”, “rotated
 985 patient position”, “limited view due to overlying hardware”). These limitations help qualify
 986 the certainty or completeness of the report’s conclusions.

987 A.2 TASK-SPECIFIC VOCABULARY CONSTRUCTION

988 To systematically capture the range of descriptive, temporal, spatial, and contextual attributes in
 989 radiologic reporting, we constructed a structured vocabulary of relation terms grounded in all schema-
 990 defined relation types instantiated in LUNGUAGE. The process followed four stages: (1) automatic
 991 candidate extraction, (2) expert review and refinement, (3) hierarchical organization into clinically
 992 meaningful subcategories, and (4) normalization of lexical variants. This pipeline was designed to
 993 maximize coverage while ensuring clinical interpretability and internal consistency.

994 **Candidate extraction.** We first piloted schema and prompt designs on 100 sample reports, iteratively
 995 refining them before applying the finalized schema to the full set of 1,473 reports (see Appendix A.3
 996 for details). Using GPT-4 (Achiam et al. (2023)), we produced initial structured outputs and extracted
 997 candidate terms corresponding to each relation type. This step emphasized high recall to capture the
 998 breadth of linguistic variation present in free-text radiology reports and provided a basis for analyzing
 999 hierarchical consistency across categories.

1000 **Expert review and refinement.** Four board-certified physicians independently reviewed the candidate
 1001 vocabularies for each relation category, verifying accurate categorization and eliminating spurious or
 1002 ambiguous expressions. Disagreements were adjudicated through consensus meetings (Figure A.6),
 1003 prioritizing clinical interpretability and reproducibility. This process was especially important
 1004 for borderline cases such as distinguishing between Condition terms under MORPHOLOGY and
 1005 subtle gradations of SEVERITY, or between field-of-view limitations and patient-related
 1006 limitations.

1007 **Hierarchical organization of entity, location, and attribute taxonomies.** All vocabularies were
 1008 organized hierarchically to reflect radiologic conventions and enable reasoning across different levels
 1009 of granularity. A comprehensive overview of the entity taxonomy (Figure A.2), location taxonomy
 1010 (Figure A.3), and attribute taxonomy (Figure A.4), together with representative examples, is provided
 1011 in Table A.1. They can be grouped into three major taxonomies:

1012 • **Entity taxonomy.** Entities were first assigned to one of six mutually exclusive Cat labels:
 1013 PF (Perceptual Findings), CF (Contextual Findings), COF (Clinical Objective Findings),
 1014 NCD (Non-CXR Diagnosis), OTH (Other Objects), and PATIENT INFO (Patient Information).
 1015 Within each label, entities were further classified into subcategories such as *Diagnostic
 1016 Observations, Anatomical Entities, Diseases and Disorders, Medical Devices, or Symptoms
 1017 & Signs*. Representative examples include: “opacity” and “right hilum” (PF), “pneumonia”
 1018 and “congestive heart failure” (CF), “oxygen saturation” (COF), “stroke” (NCD), “central
 1019 venous catheter” (OTH), and “fever” or “chronic dyspnea” (PATIENT INFO). Normalization
 1020 ensured consistent representation, while diverse raw expressions were linked at the lowest
 1021 level (e.g., “pneumonia” → “PNA,” “pneumonias”).

1022 • **Location taxonomy.** The most extensive vocabulary, comprising 546 terms, was organized
 1023 into hierarchical paths that mirror clinical localization practices. High-level systems included

1026 *respiratory* (229), *musculoskeletal* (84), *cardiovascular* (73), and *others* (160). Examples of
 1027 hierarchical paths include: “lung → lobe → right → upper,” “heart → chamber → atrium
 1028 → left,” “spine → thoracic → vertebra → T4.” This structuring enables reasoning from
 1029 coarse system-level interpretation to fine-grained anatomical localization.

1030 **• Attribute taxonomy.** Attributes were systematically organized into descriptive and temporal
 1031 axes. MORPHOLOGY (205) was divided into *shape and structure*, *texture and density*, and
 1032 *condition*. Temporal change included ONSET (57), IMPROVED (118), WORSENED (102),
 1033 and NO CHANGE (138), each stratified into graded interpretations (e.g., “moderate improve-
 1034 ment,” “minimal worsening”). Device-related metadata were captured under PLACEMENT
 1035 (74), describing both positional accuracy (e.g., “malpositioned”) and procedural changes
 1036 (e.g., “removed,” “repositioned”). Additional axes included MEASUREMENT (139), SEVER-
 1037 ITY (86), DISTRIBUTION (37), and COMPARISON (44). Auxiliary types captured contextual
 1038 but clinically relevant information: ASSESSMENT LIMITATIONS (233; e.g., “rotated pa-
 1039 tient,” “poor inspiration”), OTHER SOURCE (55; e.g., CT, MRI), and PAST HX (39; e.g.,
 1040 “status post,” “history of malignancy”). Our vocabulary was restricted to relation types that
 1041 correspond to lexically explicit attributes. Four relation types—EVIDENCE, ASSOCIATE,
 1042 DXSTATUS, and DxCertainty—were excluded. These relations are critical to the an-
 1043 notation schema but represent pragmatic inference rather than explicit lexical expressions.
 1044 For instance, EVIDENCE and ASSOCIATE encode reasoning links between entities, often
 1045 spanning sentences, while DXSTATUS and DxCertainty capture interpretive stance (e.g.,
 1046 presence vs. absence, tentative vs. definitive).

1047 **Normalization.** The resulting vocabulary includes 14 relation types derived from lexical evidence,
 1048 each normalized to a preferred set of terms and organized into semantically coherent subcategories.
 1049 We additionally performed UMLS mapping wherever possible to align relation terms with existing
 1050 biomedical ontologies, while preserving terms that fall outside conventional coverage. This ensured
 1051 both lexical consistency and clinical validity, supporting future integration. Beyond its role in structur-
 1052 ing chest X-ray reports, this vocabulary provides a reusable lexicon for tasks such as query expansion,
 1053 ontology alignment, multimodal grounding, and patient-level reasoning, thereby establishing a clin-
 1054 ically grounded and internally consistent taxonomy of radiologic language. Detailed construction
 1055 procedures are explained in Appendix A.3.

1056 **Comparison with prior resources and applications.** Compared to prior resources such as Rad-
 1057 Graph (Jain et al. (2021)), our vocabulary introduces a substantially more fine-grained taxonomy.
 1058 RadGraph defines only two entity types—*Anatomy* and *Observation*—and represents descriptive
 1059 information indirectly through coarse relations such as *modify* or *suggestive of*. In contrast, our
 1060 schema explicitly differentiates attributes such as MORPHOLOGY into *shape and structure*, *texture*
 1061 and *density*, and *condition*, and provides graded subtypes for both temporal progression and severity.
 1062 This level of granularity better reflects the linguistic practices of radiologists and enables more
 1063 nuanced downstream evaluation.

1064
 1065
 1066
 1067
 1068
 1069
 1070
 1071
 1072
 1073
 1074
 1075
 1076
 1077
 1078
 1079

Table A.1: Vocabulary Overview Taxonomy

Entity Categories				
Category	Subcategory	Example terms		
pf (Perceptual Findings)	Diagnostic Observations, Anatomical Entities, Diseases and Disorders	opacity, right hilum		
cf (Contextual Findings)	Diseases and Disorders, Diagnostic Observations	congestive heart failure, pneumonia		
cof (Clinical Objective Findings)	Diagnostic Observations, Diseases and Disorders	oxygen saturation, anti pd1 antibody		
ncd (Non-CXR Diagnosis)	Diseases and Disorders, Diagnostic Observations	stroke, seizure disorder		
oth (Other Objects)	Medical Devices, Procedures & Surgeries, Treatment & Medications	central venous catheter, lobectomy		
patient info	Symptoms & Signs, Diseases & Disorders, Treatment & Medications, Procedures & Surgeries	fever, cough, chronic dyspnea		
Attribute Categories: Spatial and Descriptive				
Category	Subcategory	Example terms		
severity	Extreme, Significant, Moderate, Mild, Minimal	moderate, severe		
measurement	Size, Quantity, Normality	2.5 cm, multiple		
morphology	Shape & Structure, Texture & Density, Condition	nodular, reticular		
distribution	Pattern, Extent, General Description	diffuse, focal		
comparison	Location & Laterality, Degree & Description	left greater than right		
Attribute Categories: Temporal Change				
Category	Subcategory	Example terms		
onset	Acute/Sudden, Chronic/Long-term, Progressive	acute, chronic		
improved	Extreme, Significant, Moderate, Mild, Minimal	resolved, decreased		
worsened	Extreme, Significant, Moderate, Mild, Minimal	enlarging, increased		
no change	No Change, Minimal Change	unchanged, persistent		
placement	Standard Position, Repositioning, New Placement, Removal, Nonstandard Position	inserted, malpositioned		
Attribute Categories: Contextual Information				
Category	Subcategory	Example terms		
assessment limitations	Evaluation, Field-of-View, Patient-Related, Technical	poor inspiration, rotated patient		
other source	Image, Signal, External Source	CT, MRI		
past hx	Past Hx	status post, known		
Location Taxonomy and Coverage				
Top-level Category	Category distribution (%)	Example Anatomical Sites	Max Depth	Example Location Paths
Respiratory	≈42%	Lungs, pleura, bronchi, thoracic wall	up to 7	<i>lung > lobes > right > upper, pleura > left > upper</i>
Cardiovascular	≈13%	Heart chambers & valves, aorta, vena cava, jugular/supra-cardiac veins	up to 6	<i>vessels > aorta > arch, heart > chambers > atrium > right, veins > jugular > internal > right</i>
Musculoskeletal	≈15%	Spine (cervical—lumbar), ribs, clavicle, shoulder & acromioclavicular joints	up to 6	<i>spine > thoracic, bones > ribs > left, joints > shoulder > right</i>
Abdominal	≈6%	Stomach, bowel segments, abdominal quadrants, sub-diaphragmatic spaces	up to 6	<i>stomach > fundus, quadrants > right, organs > intestines > duodenum</i>
Mediastinum	≈4%	Paratracheal, carinal, paramediastinal compartments	up to 5	<i>paratracheal > right, paramediastinal_region > right, carina</i>
Other structures / Descriptors	≈19%	Axilla, neck, extremities, directional descriptors, device placements	up to 5	<i>axilla > left, neck > lower, medical_device</i>

1134

1135

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

1152

1153

1154

1155

1156

1157

1158

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

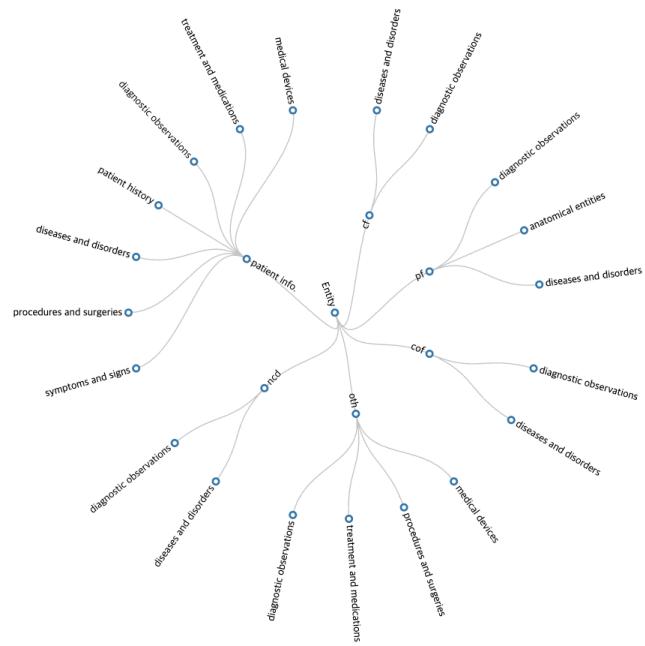
1183

1184

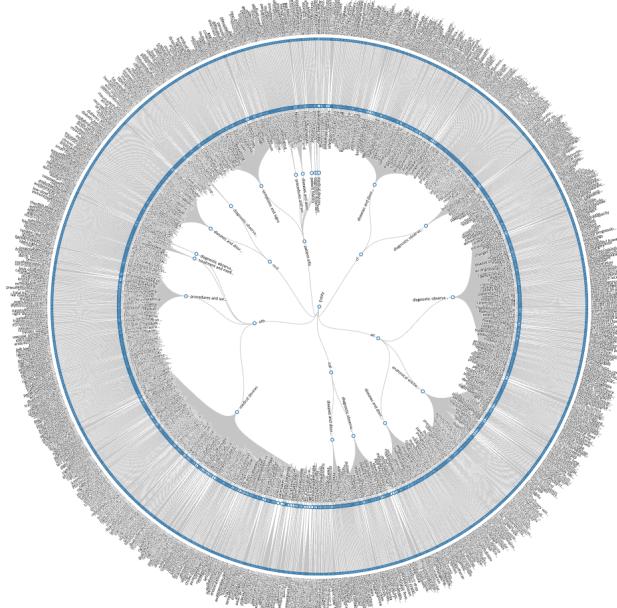
1185

1186

1187



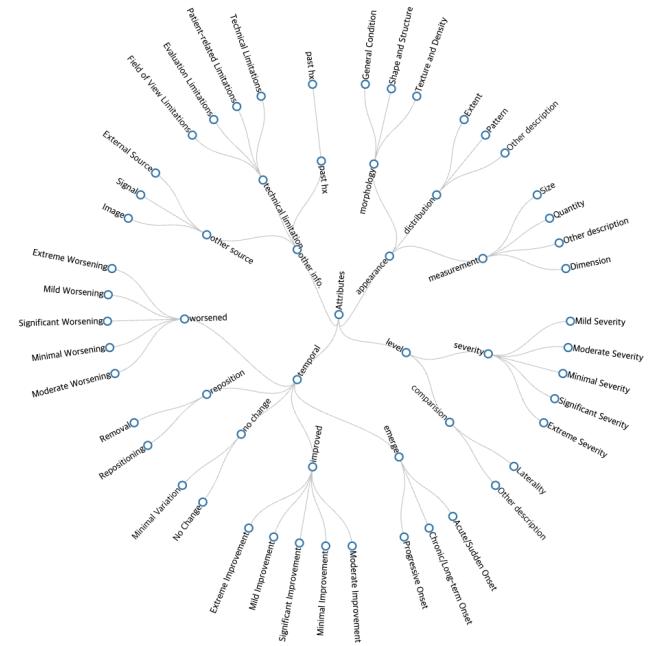
(a) Entity taxonomy (simplified). Shows only *Category* (PF, CF, OTH, COF, NCD, PATIENT INFO) and *Subcategory* (e.g., Diagnostic Observation, Anatomical Entity, Disease and Disorder).



(b) Entity taxonomy (full). Extends the simplified view by adding *Normalized terms* (canonical forms) and *Raw terms* (report expressions). For example, “pneumonia” (PF, Diagnostic Observation) is normalized to a standard form and may appear in reports as “PNA” or “pneumonia.”

Figure A.2: **Entity taxonomy.** Comparison between simplified and full versions. The simplified taxonomy shows only up to Subcategory, while the full taxonomy additionally captures normalized terms and raw report expressions.

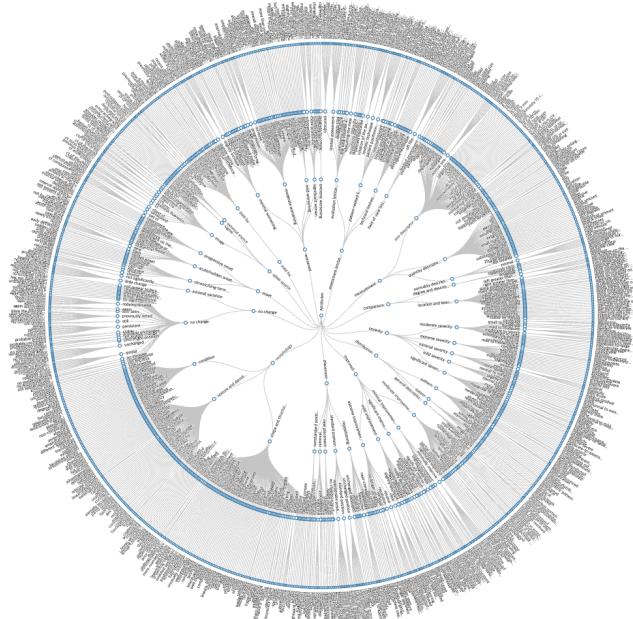
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262



1263 (a) Attribute taxonomy (simplified). Shows only *Category* (e.g.,
1264 Severity, Morphology, Distribution, Temporal change, Contextual
1265 information) and their *Subcategories* (e.g., “Extreme–Minimal”
1266 scale for Severity, “Acute/Chronic” for Onset).

1287 (b) Attribute taxonomy (full). Extends the simplified view by adding
1288 *Normalized terms* and *Raw report terms*. For example, the category
1289 *Improved* → subcategory *Minimal improvement* has normalized
1290 terms like “minimally improve” and maps to diverse raw expres-
1291 sions such as “somewhat better,” “somewhat improved,” or “slightly
1292 improved.”

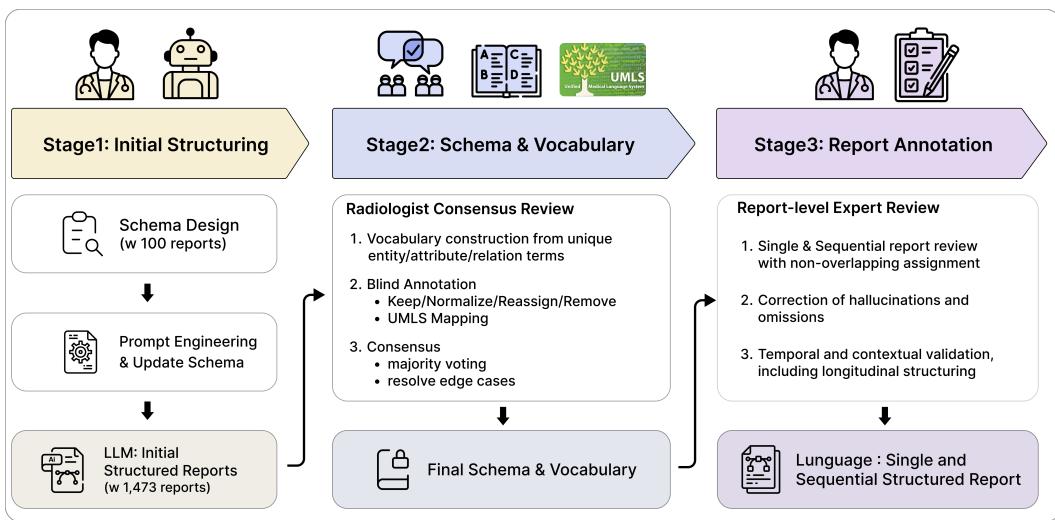
1293 **Figure A.4: Attribute taxonomy.** Comparison between simplified and full versions. The simplified
 1294 taxonomy presents categories and subcategories only, while the full taxonomy systematically incor-
 1295 porates normalized forms and raw report expressions used in clinical texts.



1287 (b) Attribute taxonomy (full). Extends the simplified view by adding
1288 *Normalized terms* and *Raw report terms*. For example, the category
1289 *Improved* → subcategory *Minimal improvement* has normalized
1290 terms like “minimally improve” and maps to diverse raw expres-
1291 sions such as “somewhat better,” “somewhat improved,” or “slightly
1292 improved.”

1296 A.3 ANNOTATION PROCESS
1297

1298 This section outlines the protocol used to ensure consistency, clinical accuracy, and transparency in our
1299 structured annotation process (Figure A.5). A key principle of our design was to adopt a *vocabulary*
1300 *first approach*. Instead of starting directly with report-level annotations, we first extracted all unique
1301 candidate terms from the structured outputs and asked radiologists to evaluate them in isolation. Each
1302 annotator independently assessed whether a term was correctly categorized under the schema (keep,
1303 normalize, reassign, or remove) and proposed refinements when necessary. This strategy served two
1304 purposes: (i) it ensured that the schema categories were well defined and clinically meaningful before
1305 applying them at scale, and (ii) it minimized inconsistencies during report annotation by locking a
1306 shared vocabulary in advance. Consensus was reached through majority voting and resolution of
1307 edge cases, producing a stable schema–vocabulary foundation upon which high-quality single and
1308 sequential report annotations were later built.
1309



1326 Figure A.5: **Annotation protocol**. The process begins with **Stage1: Initial Structuring**, where
1327 radiologists design an initial schema based on 100 chest X-ray reports and iteratively update it
1328 through prompt engineering. A large language model (LLM) then generates preliminary structured
1329 outputs from 1,473 reports, from which candidate vocabulary terms are extracted. In **Stage2: Schema**
1330 & **Vocabulary**, four board-certified radiologists conduct blinded annotation on all extracted terms
1331 (keep, normalize, reassign, or remove), supplemented with UMLS term mapping for interoperability.
1332 Consensus is reached through majority voting and resolution of edge cases, after which the schema
1333 and vocabulary are refined and locked. In **Stage3: Report Annotation**, radiologists perform
1334 expert review of non-overlapping subsets of single and sequential reports. This step includes
1335 correction of hallucinations and omissions, validation of temporal and contextual dependencies
1336 (including longitudinal structuring), and linking of cross-sentence relations such as ASSOCIATE and
1337 EVIDENCE. The pipeline yields a final schema and vocabulary as well as high-quality single and
1338 sequential structured reports (LUNGUAGE benchmark).
1339

1340 **Schema and Vocabulary Development and Validation.** The construction of the LUNGUAGE
1341 schema followed a vocabulary-first process that combined expert-driven refinement with large-scale
1342 automatic extraction. We began by randomly sampling 100 chest X-ray reports and manually drafting
1343 an initial schema. This draft was iteratively refined through prompt engineering: categories and
1344 relations were adjusted while repeatedly checking coverage against the same 100 reports. Once a
1345 stable structure was established, the process was scaled to all 1,473 reports using schema-guided
1346 prompts to a large language model, which produced preliminary structured drafts. These drafts were
1347 not treated as final structure reports, but instead served to systematically collect the entire lexical
1348 space of candidate terms across the ENTITY, ATTRIBUTE, and RELATION fields. This vocabulary-first
1349 approach ensured that the schema was grounded in actual report variation while also capturing rare or
unconventional descriptors.

1350	Attributes	Lemma (ori)	Lemma (Rev1)	Lemma (Rev2)	Synonym (ori)	Synonym (Reviewer1)	Synonym (Reviewer2)	Opinion (Reviewer1)	Opinion (Reviewer2)
1351	morphology	normal	normal	normal	normal	normal	normal	platelike add	platelike, linearly, linearly o
1352	morphology	linear	linear	linear	linear	linear, platelike	linear, linearly, linearly oriented, plate	delete - distribution(move)	
1353	morphology	patchy		patchy	patchy	patchy	patchy	nodelete	node - entity(move)
1354	morphology	nodular	nodular	nodular	nodular	nodular	nodular		calcification to (entity)
1355	morphology	calcify	calcify	calcify	calcified, calcification	calcified, calcification	calcified		
1356	morphology	streaky	streaky	streaky	streak, streaky, streaks	streak, streaky, streaks	streak, streaky, streaks		
1357	morphology	tortuous	tortuous	tortuous	tortuous, tortuosity	tortuous, tortuosity	tortuous, tortuosity		
1358	morphology	stable	stable	stable	stable	stable, unchanged, constant, gi	stable	unchanged, constant, grossly stable add	
1359	morphology	round	round	round	round, rounded	round, rounded	round, rounded		
1360	morphology	heterogeneous	heterogeneous	heterogeneous	heterogeneous, heterogeno	heterogeneous, heterogenous, heterogenous	heterogeneous, heterogenous, heterogenous	inhomogeneous add	
1361	morphology	loculated	loculated	loculated	loculated	loculated	loculated		
1362	morphology	hazy	hazy	hazy	hazy	hazy	hazy		
1363	morphology	dense	dense	dense	denser, dense	denser, dense	denser, dense, densely		
1364	morphology	prominent	prominent	prominent	prominent, prominence	prominent, prominence, frank,	prominent, prominence	frank, exuberant, marked add	
1365	morphology	widen	widen	widen	widened	widened, wide, extensive	widened	wide, extensive add	
1366	morphology	enlargement	enlargement	enlargement	enlarged, enlargement, dilat	enlarged, enlargement, dilat	enlarged, enlargement		dilated separate morpholo
1367	morphology	faint	faint	faint	faint, vague	faint, vague	faint, vague		
1368	morphology	coil	coil	coil	coiled, looped	coiled, looped, curled up, curin	coiled, looped	curled up, curling, curled, kinking add	
1369	morphology	residual	residual	residual	residual	residual	residual	delete - no change(move)	delete - location(move)
1370	morphology	peribronchial		peribronchial				delete - distribution(move)	delete - location(move)

Figure A.6: **Example of consensus protocol for vocabulary validation.** This figure illustrates the vocabulary validation process using the *morphology* category as an example. For illustration, only the results from two reviewers are shown, although in the real protocol, all four radiologists independently reviewed every candidate term. The column *Lemma (ori)* represents raw vocabulary extracted from the LLM. Reviewers evaluated each term to determine whether it should remain in the category, be reassigned, merged as a synonym, or removed, and also identified appropriate synonyms. Green highlights indicate terms where reviewer opinions diverged, requiring consensus discussion. While not shown in this example, the full protocol also incorporated UMLS mappings and subcategory assignments.

The candidate vocabulary was then subjected to blinded review by four board-certified radiologists (an example is shown in Figure A.6). For each term, annotators determined whether to *keep*, *normalize*, *reassign*, or *remove*, and assigned it to an appropriate schema category (e.g., PF, CF, OTH). Disagreements were resolved through majority voting, while ambiguous cases were recorded in an *edge-case log* and revisited during consensus meetings. To promote alignment with established clinical practice and interoperability, raw terms were cross-checked against Fleischner Society terminology (Bankier et al. (2024)), particularly during normalization, and mapped to UMLS concepts when appropriate (e.g., UMLS TERM (CODE: C12345)).

UMLS API

Enter your search term

pneumonia

Select the relation type

children

Search

	name	ui	rootSource
0	Pneumonia	C0032285	MTH
1	Streptococcal pneumonia	C0155862	MTH
2	Aspiration Pneumonia	C0032290	MTH
3	Mycoplasma pneumonia	C0032302	MTH
4	Pneumonia, Viral	C0032310	MSH

Figure A.7: **UMLS mapping tool.** A custom interface was developed to retrieve candidate concepts via the UMLS API. Radiologists manually reviewed suggested mappings to select the most semantically aligned concepts, while terms without clear matches were explicitly marked as unmapped (-).

We developed a custom UMLS mapping tool (Figure A.7) to retrieve candidate concepts via the UMLS API, and radiologists manually reviewed these candidates to select the most semantically

aligned concepts. Terms without a clear correspondence were not forced into external standards; instead, they were explicitly marked as unmapped (– (CODE: –)). All mappings, including unmapped entries, were preserved in the final vocabulary for transparency and reuse.

Through this iterative adjudication and validation process, the vocabulary evolved from a raw set of extracted terms into a clinically coherent schema encoding core entity categories, relation types (e.g., DXCERTAINTY, ASSOCIATE, EVIDENCE), and detailed attributes (e.g., MORPHOLOGY, MEASUREMENT, TECHNICAL LIMITATION). The outcome was a locked schema and vocabulary that balanced comprehensive coverage with clinical precision, establishing the foundation for all subsequent annotation in the LUNGUAGE benchmark.

Report Annotation and Validation. After the schema and vocabulary were finalized, structured report annotation proceeded under a standardized workflow designed to ensure both consistency and clinical reliability. Radiologists worked with shared resources, including (i) a finalized schema document defining all entity, relation, and attribute types and their origin (image-derived vs. context-derived), (ii) an evolving *edge-case log* recording ambiguous examples and their resolutions, and (iii) a set of standardized decision rules for term normalization and schema assignment.

Using a custom annotation interface (Figure A.8), annotators reviewed non-overlapping subsets of reports to avoid redundancy. Their responsibilities extended beyond sentence-level checks to include linking observations across sentences (e.g., ASSOCIATE, EVIDENCE), validating temporal and contextual dependencies, and identifying hallucinations or omissions introduced during LLM-based structuring. This two-stage pipeline—automatic structuring followed by expert curation—yielded high-quality annotations for all 1,473 single reports and 186 longitudinal cases over a six-month period, with radiologists contributing multiple hours per week and participated in weekly review meetings. The resulting resource demonstrates annotation consistency and clinical reliability, providing a validated foundation for the LUNGUAGE benchmark.

A.3.1 SINGLE REPORT ANNOTATION DETAILS

Single Structured Report Statistics

- Total number of reports: 1,473 chest X-ray reports
- Total number of patients: 230
- Number of imaging studies per patient: Ranges from 1 to 15
- Total number of annotated entities: 17,949
- Total number of annotated relation–attribute pairs: 23,307

To construct a clinically reliable gold-standard dataset, we implemented a structured annotation pipeline that reviewed and refined the initial triplets (entity-relation-attribute) generated by GPT-4 (0613). Unlike the vocabulary construction phase—which focused on individual terms without considering report context—this stage involved section-by-section review of all structured outputs in each report to ensure contextual accuracy and logical consistency.

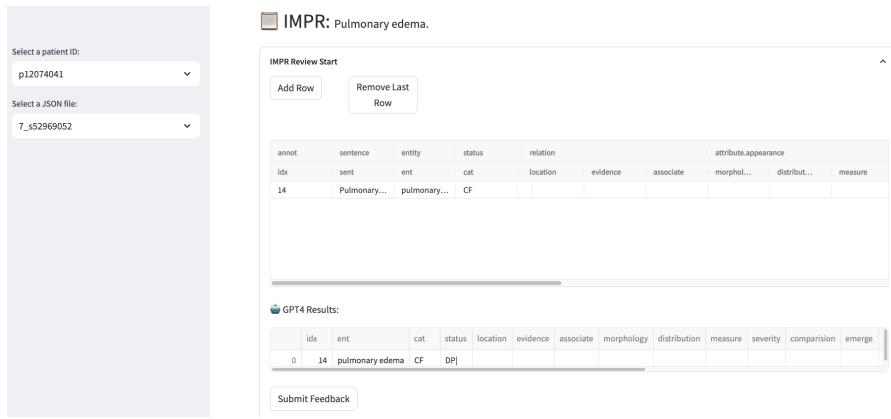
All 1,473 chest X-ray reports in LUNGUAGE were divided evenly among annotators. Each annotator independently reviewed approximately one-quarter of the dataset, ensuring balanced coverage and minimizing reviewer bias across the annotated corpus. Within each report, annotators examined the structured outputs across the *history/indication*, *findings*, and *impression* sections. The goal was to verify whether the extracted (*entity, relation, attribute*) triplets accurately captured the meaning of the source text and aligned with the predefined schema.

This review explicitly included schema elements that require contextual interpretation and cannot be evaluated at the lexical level alone—namely, DXSTATUS, DXCERTAINTY, ASSOCIATE, and EVIDENCE. These attributes reflect interpretive judgments, such as identifying when an “opacity” supports a diagnosis of “pneumonia” or whether two entities should be linked through an associative relation. Annotators verified whether such relations were correctly inferred from the surrounding text and whether the attributes assigned to each entity (e.g., presence, uncertainty, temporal change) matched the narrative context.

To support this process, we developed a custom annotation interface (Figure A.8) that displayed the original report text alongside GPT-4’s predicted triplets and an editable table of structured fields.

1458 Each sentence in the report was paired with its associated annotations, including entity category,
 1459 relation type, and all relevant attributes. Annotators could directly add, edit, remove, or merge entries
 1460 to reflect clinically accurate interpretations. For example, terms like “ground glass opacity”—which
 1461 could be mistakenly split—were merged into a single PF (perceptual finding) entity based on
 1462 how radiologists commonly use the phrase. Annotation was conducted separately for each section
 1463 (history, findings, impression), and the interface supported sentence-level review within each
 1464 section to ensure consistent entity–relation mappings when terms appeared across multiple sentences.

1465 As a result of this process, the finalized gold-standard dataset includes 17,949 validated entities
 1466 and 23,307 relation instances. These annotations encompass both explicit descriptive attributes
 1467 and contextually inferred diagnostic relationships, providing a robust benchmark for evaluating
 1468 schema-based information extraction systems in chest radiograph interpretation. As an illustration,
 1469 Figure A.9 shows the knowledge graph representation of a single annotated report drawn from the
 1470 dataset.



1485 Figure A.8: Annotation interface used during gold dataset construction. Annotators reviewed GPT-
 1486 4-generated triplets per report section and refined the entity–relation structure to ensure schema
 1487 correctness and contextual validity.

1488
 1489
 1490
 1491
 1492
 1493
 1494
 1495
 1496
 1497
 1498
 1499
 1500
 1501
 1502
 1503
 1504
 1505
 1506
 1507
 1508
 1509
 1510
 1511

1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565

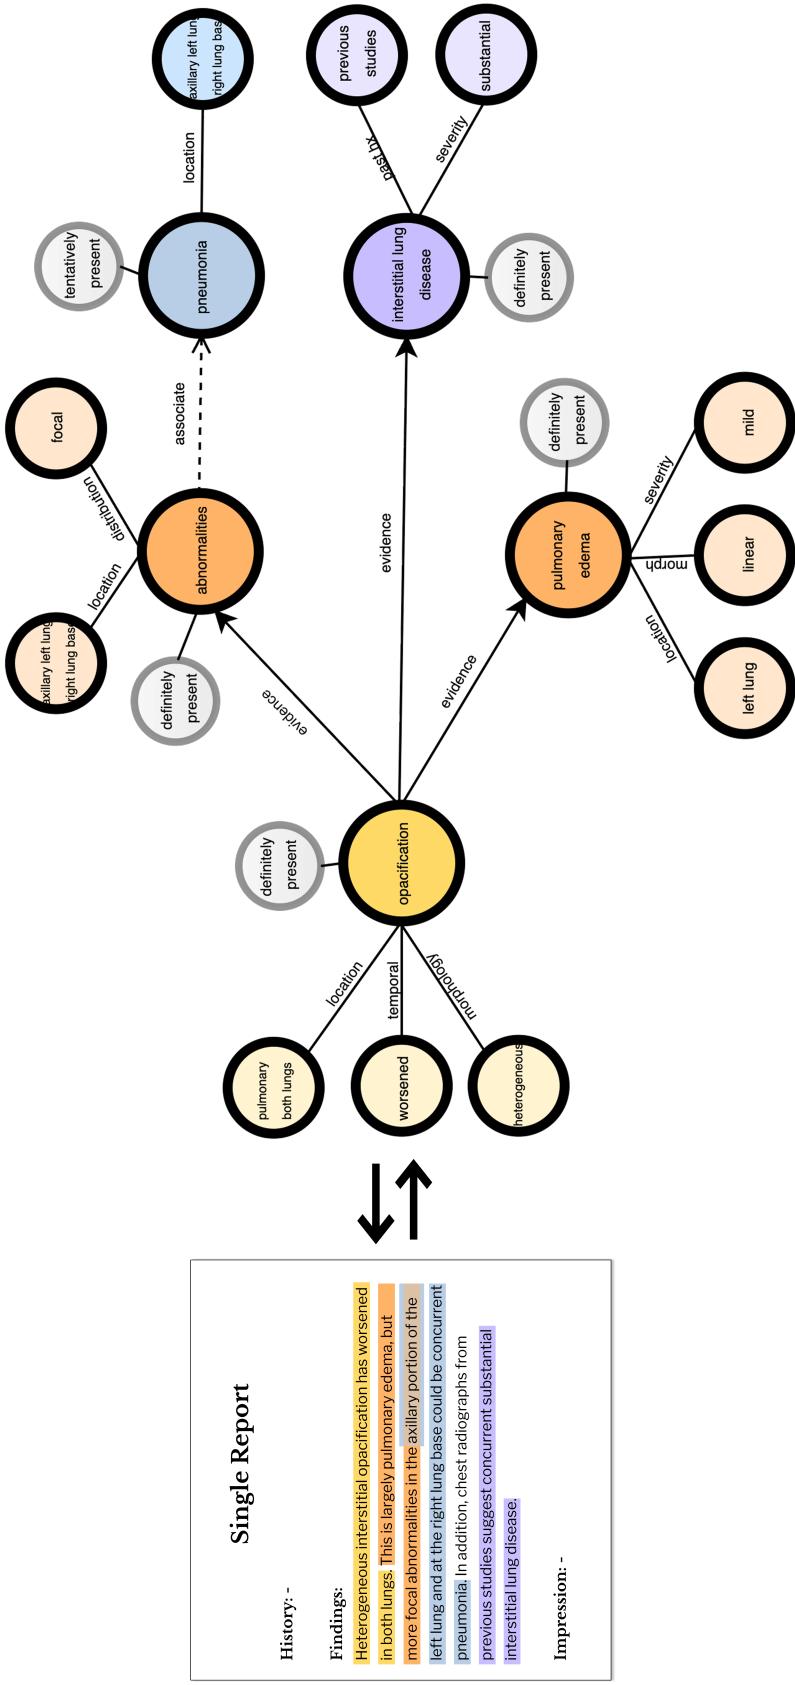


Figure A.9: **Structured Report as Knowledge Graph**. The figure illustrates how the findings section of a single chest X-ray report (bottom box) is structured under our schema. The text is decomposed into entities and their relations, including entity–entity links (e.g., associate, evidence) and entity–attribute links (e.g., status, certainty, location, severity, morphology). Connecting these components yields a knowledge graph (KG) that makes explicit both the relationships among entities and the descriptive properties attached to each entity.

1566
1567

A.3.2 SEQUENTIAL REPORT ANNOTATION DETAILS

1568
1569

Sequential Structured Report Statistics

1570
1571

- Total number of reports: 186 chest X-ray reports
- Total number of patients: 30 (subset of the 230-patient cohort)
- Reports per patient: between 2 and 14
- Time intervals between reports: from 1 day to 1,200 days
- Observation pairs: 95,404 in total. This number comes from comparing every possible pair of annotated entities (observations) within each patient trajectory (i.e., all $\binom{n}{2}$ combinations). For example, a patient trajectory containing 34 entities yields 561 pairwise comparisons, while a dense case with 141 entities results in 9,870 pairs.

1572
15731574
15751576
1577

- Total number of reports: 186 chest X-ray reports
- Total number of patients: 30 (subset of the 230-patient cohort)
- Reports per patient: between 2 and 14
- Time intervals between reports: from 1 day to 1,200 days
- Observation pairs: 95,404 in total. This number comes from comparing every possible pair of annotated entities (observations) within each patient trajectory (i.e., all $\binom{n}{2}$ combinations). For example, a patient trajectory containing 34 entities yields 561 pairwise comparisons, while a dense case with 141 entities results in 9,870 pairs.

1578
1579

In contrast to the single-report structuring phase, which focused on refining schema-based annotations within individual reports, the sequential annotation phase aimed to assess the longitudinal consistency of entity-level interpretations across temporally ordered reports from the same patient. This required global comparisons across all sections—history, findings, and impression—integrating entity–relation triplets into clinically coherent sequences.

1580
1581

Unlike earlier phases that processed each report independently, this step involved exhaustive pairwise comparisons of all annotated expressions across time. Annotators judged whether lexically distinct phrases referred to the same underlying clinical entity by examining radiological terminology, anatomical location, temporal modifiers (e.g., “resolving”, “unchanged”), and diagnostic specificity. Expressions identified as referring to the same finding were grouped together; otherwise, they were assigned to separate entity groups.

1582
1583

To further structure these entity groups, we assessed whether each represented a single episode of care or multiple distinct episodes. This required examining the temporal order and interval between observations. Intervals were computed using the StudyDate metadata from MIMIC-CXR, and episode boundaries were assigned based on temporal coherence—considering factors such as time gaps, patterns of resolution or worsening, and recurrence of findings.

1584
1585

For example, a progression from “moderate left effusion” (day 0) to “small effusion” (day 14) and “trace effusion” (day 45) was treated as a single resolving episode. However, a subsequent “moderate effusion” on day 180 was regarded as a separate episode, while all entities assigned to either episode are grouped into the same Entity Group. Similarly, “right lower lobe opacity” followed by “resolving infiltrate” was interpreted as one episode, whereas a new “opacity” on day 150 initiated a different episode. This process was applied to 186 chest X-ray reports from 30 patients, yielding longitudinal annotations that capture consistent entity grouping across lexical variations and clinically coherent organization of episodes based on temporal reasoning.

1586
1587

To better characterize the annotation results, we summarize the distribution of entity groupings and temporal episodes in Table A.2. The columns report:

1588
1589

- **# Reports:** The total number of reports per patient sequence.
- **Entity Group Distribution:** The number of findings assigned to each entity group (#Group), after normalization and longitudinal reasoning. Some groups consist of a single unique expression, while others aggregate multiple semantically related terms.
- **Temporal Group Distribution:** The number of findings assigned to each temporal group (#Group), where each group represents a distinct clinical episode.

1620 Table A.2: Distribution of entity groups and temporal groups across annotated patient sequences.
1621

1622 Subject ID	1623 # Reports	1624 Entity Group Distribution (#Group:Count)	1625 Temporal Group Distribution (#Group:Count)
1626 p10046166	1627 6	1:26, 2:3, 3:3, 5:1, 7:1	1:32, 2:2
1628 p10274145	1629 5	1:19, 2:11, 3:2, 4:3	1:33, 2:2
1630 p10523725	1631 9	1:36, 2:6, 3:3, 4:2, 5:2, 7:2	1:47, 2:2, 3:1, 6:1
1632 p10532326	1633 5	1:36, 2:6, 6:1	1:42, 2:1
1634 p10885696	1635 8	1:33, 2:9, 3:5, 5:1, 6:2, 12:1	1:45, 2:2, 3:4
1636 p10886362	1637 10	1:26, 2:3, 3:6, 4:4, 6:1, 7:1, 9:1, 13:1	1:39, 2:4
1638 p10959054	1639 7	1:31, 2:6, 3:2, 4:2, 5:1, 6:1, 9:1	1:37, 2:5, 3:2
1640 p11540283	1641 5	1:26, 2:5, 3:1, 4:2	1:33, 4:1
1642 p11607628	1643 8	1:11, 2:2, 3:2, 4:4, 5:1, 6:2, 7:1, 8:1, 10:1	1:23, 2:2
1644 p11879886	1645 6	1:27, 2:10, 3:3, 4:3, 5:2, 6:1	1:39, 2:4, 3:2, 4:1
1646 p12433421	1647 13	1:49, 2:6, 3:10, 5:1, 7:1, 17:1	1:66, 2:2
1648 p12966004	1649 3	1:21, 2:10, 4:1, 5:1	1:27, 2:5, 3:1
1650 p15094735	1651 2	1:8, 2:5, 3:2, 4:1	1:13, 2:3
1652 p15109122	1653 4	1:11, 2:1, 3:4, 4:1	1:16, 2:1
1654 p15207316	1655 4	1:16, 2:6, 3:4	1:26
1656 p15272972	1657 5	1:10, 2:3, 3:3, 4:2, 5:1	1:17, 2:2
1658 p15321868	1659 6	1:24, 2:5, 3:2, 4:1, 5:2	1:32, 2:2
1660 p15446959	1661 5	1:29, 2:7, 3:3, 4:2	1:37, 2:4
1662 p15881535	1663 3	1:17, 2:2, 3:2, 5:1	1:20, 2:2
1664 p16059470	1665 4	1:29, 2:5, 3:3, 6:1	1:37, 2:1
1666 p17270742	1667 5	1:25, 2:4, 3:5, 4:2, 5:1	1:30, 2:4, 3:3
1668 p17288844	1669 6	1:42, 2:8, 3:2, 4:1, 5:1	1:54
1670 p17396677	1671 4	1:18, 2:3, 3:3, 4:1	1:23, 2:2
1672 p17720924	1673 8	1:30, 2:8, 3:5, 4:1, 5:1	1:41, 2:2, 4:2
1674 p17962324	1675 5	1:37, 2:4, 3:3, 4:1	1:43, 2:1, 3:1
1676 p18079481	1677 14	1:34, 2:10, 3:3, 4:2, 6:3, 7:3, 8:1	1:43, 2:10, 3:3
1678 p18417750	1679 7	1:41, 2:10, 3:5, 9:1	1:47, 2:7, 3:1, 6:2
1680 p18517718	1681 6	1:15, 2:2, 3:4, 4:2, 5:1, 7:1	1:25
1682 p18570152	1683 5	1:23, 2:4, 3:2, 4:1, 5:1, 6:2	1:25, 2:4, 3:3, 4:1
1684 p19150427	1685 8	1:42, 2:7, 3:2, 4:1, 6:1	1:49, 2:2, 3:1, 4:1

1646 Across the 30 patients in the sequential evaluation phase, the number of temporal groups assigned
 1647 to a single entity group ranged from 1 to 6, indicating that some findings were observed in multiple
 1648 distinct clinical episodes over time. Likewise, the number of distinct entity groups varied significantly.
 1649 Most entity groups consisted of a single mention, but some aggregated up to 17 lexically different
 1650 expressions. For example, subject p12433421 exhibited the most diverse entity grouping, with 17
 1651 distinct phrases all referring to variations of pleural effusion (e.g., “effusion,” “pleural effusion,”
 1652 “pleural effusion left”) unified under one normalized cluster. Similarly, subjects p10523725 and
 1653 p18417750 exhibited high temporal discontinuity, with single entity groups spanning up to 6 distinct
 1654 episodes (e.g., recurrent dyspnea separated by periods of resolution). These results highlight the
 1655 complexity and variability of radiologic expression in longitudinal reporting, and underscore the
 1656 necessity of models and metrics capable of robustly handling both semantic variation and episodic
 1657 continuity in time-aware clinical tasks.

1658
 1659
 1660
 1661
 1662
 1663
 1664
 1665
 1666
 1667
 1668
 1669
 1670
 1671
 1672
 1673

1674 A.4 DETAILED COMPARISON WITH EXISTING STRUCTURED-REPORT DATASETS
1675

1676 To clarify the design of our schema, we compare it against prior structured-report datasets using
1677 a common criterion. For fair comparison, all resources are re-aligned to a unified Named Entity
1678 Recognition (NER) / Relation Extraction (RE) standard: core clinical concepts are retained as **entities**,
1679 while descriptive aspects (e.g., status, measurement, severity) are recategorized as **relation labels**.
1680 This alignment can differ from the original dataset definitions. For example, RadGraph (Jain et al.
1681 (2021)) defines *Anatomy* and *Observation* as entities, but further subdivides observations into three
1682 uncertainty levels (*Definitely Present*, *Uncertain*, *Definitely Absent*), yielding four entity categories.
1683 In our comparison, however, only *Anatomy* and *Observation* are retained as entities, while the
1684 uncertainty levels are reassigned as *relation labels* linked to the observation entity. Consequently,
1685 the counts in Table A.3 may not exactly match those in the original papers, since all datasets are
1686 reorganized under a single consistent criterion to enable direct side-by-side comparison.
1687

1688 Table A.3: Comparison of schema coverage across major structured-report datasets. Numbers
1689 indicate the count of categories after re-aligning each schema under a unified NER/RE schema.
1690 **Entities** represent the number of core concept types defined in each dataset. **Entity-Entity relations**
1691 are semantic links between two entities (e.g., *located at*, *associate*, *evidence*). **Entity-Attribute**
1692 **relations** describe properties attached to a single entity (e.g., *status*, *certainty*, *location*, *severity*,
1693 *morphology*). **Sequential relation** indicates whether temporal continuity across multiple reports is
1694 explicitly modeled. Because each dataset originally adopted its own definitions, the numbers here
1695 may not match the original papers exactly.

1695 Dataset	#Entities	#Entity-Entity Relations	#Entity-Attribute Relations	Sequential Relation
1696 RadGraph	2	3	2	✗
1697 RadGraph-XL	2	3	3	✗
1698 RadGraph-2	3	3	9	✗
1699 Rate-NER	3	0	1	✗
1700 CAD-Chest	1	0	4	✗
1701 LUNGUESCORE	6	2	16	✓

1702
1703 (1) ENTITY-LEVEL DIAGNOSTIC SOURCE

1704 Existing structured-report datasets define only a limited range of entity types and do not explicitly
1705 distinguish whether a concept is directly inferable from chest radiographs. For example, RadGraph
1706 (Jain et al. (2021)) and RadGraph-XL (Delbrouck et al. (2024)) include only *anatomy* and *observation*,
1707 while RadGraph-2 (Khanna et al. (2023)) adds *device*. Rate-NER (Zhao et al. (2024)) defines *anatomy*,
1708 *abnormality*, and *disease*, whereas CAD-Chest (Zhang et al. (2023)) reduces coverage to a single
1709 entity type (*disease*).
1710

1711 This design mixes image-grounded findings (e.g., opacity, atelectasis) with contextual or diagnostic
1712 terms (e.g., pneumonia, heart failure). As a result, models may hallucinate unsupported content or be
1713 penalized unfairly when generating clinically valid but context-dependent descriptors.
1714

1715 Our schema introduces six categories with an explicit separation by visual inferability: **PF**, **CF**,
1716 and **OTH** represent image-grounded entities, while **COF**, **NCD**, and **PATIENT INFO** capture
1717 information that cannot be reliably inferred from the image itself. This distinction improves training
1718 relevance, enables fairer evaluation, and reduces hallucination risk.
1719

1720 (2) RELATION-LEVEL SEMANTIC PRECISION

1721 Relation definitions in prior datasets are generally coarse. RadGraph variants restrict entity–entity
1722 links to three types (*modify*, *located at*, *suggestive of*) and entity–attribute links to *status* and *certainty*
1723 (later including *measurement*). Rate-NER and CAD-Chest define only one to four relation categories
1724 in total.
1725

1726 Such limited taxonomies conflate distinct clinical reasoning cues. For instance, RadGraph’s *suggestive*
1727 of combines associative reasoning (e.g., “right lung opacity may be nipple shadow”) with evidential
1728 claims (e.g., “opacity suggests pneumonia”), which are clinically distinct. Temporal descriptors such

1728 as “improved,” “worsened,” or “no change” are also collapsed into generic status labels, obscuring
 1729 clinically meaningful differences.
 1730

1731 Our schema expands this design to **18 relation labels** that provide more fine-grained semantic
 1732 distinctions. This includes entity–entity relations such as *associate* and *evidence*, as well as a rich set
 1733 of entity–attribute relations covering *status*, *certainty*, *severity*, *location*, *morphology*, *measurement*,
 1734 *onset*, *improved*, *worsened*, *no change*, *placement*, *past history*, *other source*, and *assessment*
 1735 *limitation*. This expanded taxonomy enables more precise error analysis and ensures that distinct
 1736 error types (e.g., incorrect severity vs. incorrect negation) are penalized appropriately.
 1737

1738 (3) PATIENT-LEVEL LONGITUDINAL LINKAGE

1739 All prior structured-report datasets treat each report in isolation, preventing validation of temporal
 1740 descriptors or consistency across multiple studies. They cannot verify whether a reported improvement
 1741 aligns with prior findings, nor can they unify lexical variants (e.g., “opacity” vs. “consolidation”)
 1742 across timepoints.

1743 Our schema explicitly incorporates longitudinal structure through two constructs: **ENTITYGROUPS**
 1744 and **TEMPORALGROUPS**. **ENTITYGROUPS** unify lexically different mentions of the same
 1745 clinical finding across reports of a patient, ensuring consistent recognition of semantically equivalent
 1746 terms. **TEMPORALGROUPS** segment patient trajectories into clinical episodes, reflecting disease
 1747 progression, resolution, or chronic persistence.

1748 Together, these mechanisms allow evaluation to assess whether generated findings remain temporally
 1749 coherent across multiple visits, aligning with radiological practice where longitudinal comparison is
 1750 central. By introducing sequential linkage, our schema enables structured evaluation of longitudinal
 1751 reasoning—an ability absent from prior resources.
 1752

1753
 1754
 1755
 1756
 1757
 1758
 1759
 1760
 1761
 1762
 1763
 1764
 1765
 1766
 1767
 1768
 1769
 1770
 1771
 1772
 1773
 1774
 1775
 1776
 1777
 1778
 1779
 1780
 1781

1782 **B TWO-STAGE STRUCTURING FRAMEWORK DETAILS**
17831784 **B.1 SINGLE STRUCTURING PROMPT**
17851786 **Prompt Template for Single Structuring**
17871788 You are a high-precision relation-extraction engine for chest X-ray report sections.
1789 Given a structured input, extract clinical relations between entities while strictly
1790 conforming to the provided schema and labeling rules.

1791 Your task:

- Identify valid entity pairs and annotate appropriate relation types between them.
- Assign Cat, Dx_Status, and Dx_Certainty labels to subject entities.
- Use the provided "candidates" field to guide your extraction and ensure spelling/casing consistency.
- For each identified relation, include:
 - subject_ent: unique index of subject entity
 - subject_cat: entity type
 - obj_ent_idx: unique index of related object
 - relation: relation type (must be one of the allowed relations)
 - sent_idx: sentence index from which the relation is derived
- Output a JSON object that conforms to the "PyraDict StructuredOutput" schema.
- Do not return natural language commentary or raw triples.

1799 Input JSON format:

```
{
  "report_sections": [
    {
      "sent_idx": 1,
      "sentence": "Findings suggest possible pneumonia in the right lower lobe with opacity.",
      "candidates": [
        ["Pneumonia", ["Entity1"]],
        ["Findings", ["Entity1"]],
        ["Right lower lobe", ["Location1"]],
        ["Opacity", ["Entity1"]]
      ]
    },
    {
      "sent_idx": 2,
      "sentence": "A new small pleural effusion is seen on the left side.",
      "candidates": [
        ["Pleural effusion", ["Entity1"]],
        ["Left side", ["Location1"]],
        ["New", ["None"]],
        ["Small", ["Measurement1"]]
      ]
    }
  ]
}
```

1817 Allowed Relation Types:

- Cat, Status, Location, Placement, Associate, Evidence, Morphology, Distribution, Measurement, Severity, Comparison, Onset, No Change, Improved, Worsened, Past Hx, Other Source, Assessment Limitations

1820 Labeling Rules Summary:

- Every subject entity must be assigned exactly one of: Cat, Dx_Status, Dx_Certainty.
- Placement is used only for spatial position.
- Placement is disallowed for devices (Cat = OTH).
- Evidence relations must point from diagnoses to radiological findings.
- Attribute relations must be explicitly stated in the sentence.

1825 Output:

- A list of structured entries containing entity and relation annotations.
- Output must be a single valid JSON object.
- Include only entities mentioned in the text but not in the candidates, following the ent_idx order of appearance.

1829 For example:

```
Input: <related_example>
Output: <structured_report>.
```

1831

1833

1834

1835

1836 B.1.1 VOCABULARY MATCHING ALGORITHM
1837

1838 To improve consistency in entity extraction and reduce hallucinations in schema-based structuring, we
1839 implemented a vocabulary-guided span matching algorithm (see Appendix A.2 for details on vocabu-
1840 lary construction). This algorithm processes each section of the radiology report (e.g., *findings*) to
1841 identify candidate entity spans by directly matching contiguous token sequences against entries in a
1842 schema-defined vocabulary, without normalization such as lowercasing or punctuation removal. Each
1843 sentence is evaluated independently, and multiple overlapping matches are retained—e.g., “left lung”
1844 may correspond to both *PF* and *LOCATION*.

1845 Importantly, the matched vocabulary spans are not assumed to constitute a complete or authoritative
1846 set of entities. Instead, they serve as reference cues for the LLM, which remains responsible for
1847 the final relation extraction. The LLM is expected to leverage the matched terms as guidance while
1848 retaining the flexibility to identify additional entities or values not covered by the vocabulary. This
1849 design accommodates incompleteness in the vocabulary and enables the model to make context-
1850 sensitive inferences based on both the prompt and observed patterns in the data.

1851 The matching algorithm is summarized below:

1853 **Algorithm 1** Span-Based Vocabulary Matching
1854

1: **Input:** Curated vocabulary V ; report section T composed of multiple sentences.
2: **Output:** List of matched word spans in T , each labeled with one or more schema categories.
3: Build a dictionary V_{lookup} from surface forms in V , mapping each to one or more associated
schema categories.
4: **for** each sentence s in T **do**
5: Split s into a sequence of n words, each with character-level start and end offsets
6: **for** span length l from n down to 1 **do**
7: **for** start index $i = 0$ to $n - l$ **do**
8: Extract word span $s_{i:i+l}$ and its character range from original sentence
9: Query V_{lookup} for exact match of the word span
10: **if** match found **then**
11: **for** each schema category linked to the matched term **do**
12: Record span text, character start/end indices, matched term, and category
13: **end for**
14: **end if**
15: **end for**
16: **end for**
17: **end for**
18: **return** List of matched spans with associated categories

1874 This procedure constrains entity recognition to schema-aligned expressions, allowing the LLM to
1875 focus on inferring relational structure rather than determining precise span boundaries. By anchoring
1876 extraction to predefined lexical targets, it reduces ambiguity and ensures consistent treatment of
1877 clinically equivalent yet lexically variable expressions.

1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889

1944 B.3 SINGLE STRUCTURING ANALYSIS
1945
1946
1947
1948

Table B.1: Ablation results of GPT-4.1 under varying prompt-shot configurations and vocabulary matching. We report precision (P), recall (R), and F1 scores for both entity-relation pair extraction and complete triplet extraction tasks.

1949 1950 1951	1952 1953	1954 1955	1956	entity-relation			entity-relation-attribute		
				1957 1958	1959	1960	1961 1962 1963 1964	1965 1966 1967 1968 1969 1970 1971 1972	
Shot	Vocab Usage	F1	P	R	F1	P	R		
Zero	No	0.79	0.65	1.00	0.52	0.65	0.44		
	Yes	0.92	0.85	1.00	0.78	0.80	0.77		
5-shot	No	0.93	0.87	1.00	0.84	0.85	0.83		
	Yes	0.94	0.89	1.00	0.87	0.87	0.86		
10-shot	No	0.94	0.88	1.00	0.86	0.86	0.85		
	Yes	0.96	0.91	1.00	0.89	0.90	0.87		

We conducted an ablation study to quantify the individual and combined effects of vocabulary matching and in-context demonstrations on single-report structuring. Using 80 radiology reports from 30 patients, previously annotated for sequential evaluation, this subset enabled consistent evaluation across controlled input conditions.

Six configurations were tested by varying two factors: (1) whether span-to-category alignment via vocabulary matching was applied, and (2) the number of in-context examples provided in the prompt (0, 5, or 10). Vocabulary matching involved matching contiguous text spans against a predefined lexicon and retrieving all associated schema categories, ensuring lexical consistency and reducing ambiguity in span interpretation, as described in Appendix B.1.1. In-context demonstrations consisted of structured examples retrieved from the gold set of structured reports using BM25 retrieval, based on textual similarity to the input report. These examples illustrate appropriate usage of entity types and relations under the schema.

As shown in Table B.1, vocabulary matching consistently enhanced performance across all prompt configurations. Under the zero-shot setting, incorporating vocabulary guidance raised the triplet-level F1 score from 0.52 to 0.78, and the entity-relation F1 from 0.79 to 0.92. When five in-context demonstrations were provided, the triplet F1 increased further—reaching 0.84 without vocabulary and 0.87 with vocabulary. The highest accuracy was achieved by combining both components: the 10-shot setting with vocabulary matching attained a triplet F1 of 0.89.

These results indicate that vocabulary matching and in-context demonstrations offer complementary benefits. Vocabulary alignment improves lexical grounding and category consistency, while prompting with examples strengthens structural fidelity across varying linguistic expressions. Together, they establish a robust configuration for producing schema-compliant structured outputs from free-text radiology reports.

To illustrate the qualitative impact of vocabulary matching and prompt-based demonstrations, we examined example outputs across configurations with and without these components. In the sentence “*there is no focal consolidation*”, the model without vocabulary and prompt guidance extracted “*focal consolidation*” as the entity, conflating the modifier and the core clinical concept. In contrast, all other configurations correctly identified “*consolidation*” as the schema-aligned entity. A similar pattern was observed in “*there are no new focal opacities concerning for pneumonia*”, where the no-guidance setup extracted “*focal opacities*”, whereas guided configurations yielded the correct entity “*opacities*”.

These examples underscore the importance of explicitly aligning model outputs to a predefined schema. Linguistically valid but structurally inconsistent extractions can hinder downstream applications, where precise interpretation and reliable information linkage are essential. By providing lexical anchoring through vocabulary and structural demonstrations via prompts, our approach ensures that model predictions are not only accurate but also semantically coherent and clinically usable.

1998
1999

B.4 SEQUENTIAL STRUCTURING ANALYSIS

2000 We qualitatively evaluated model behavior in the sequential setting by analyzing entity grouping
2001 outputs over time. Using longitudinal chest X-ray reports from representative patients, we assessed
2002 how well the predicted entity groupings aligned with gold-standard annotations. As illustrated in
2003 Figure B.1, we examined diverse cases to understand temporal consistency and grouping granularity.2004 In general, clinical observations were consistently grouped across both annotations. For instance, in
2005 Patient p15881535, three lexical variants—*orthopedic side plate right clavicular unchanged, right*
2006 *clavicle hardware*, and *internal fixation hardware*—were all correctly assigned to the same entity
2007 group in both the gold standard and the model output. Although the representative phrase differed,
2008 the group identity was preserved, indicating successful recognition of referential equivalence across
2009 timepoints.2010 Discrepancies primarily arose from differences in granularity rather than semantic errors. In the case
2011 of Patient p15881535, temporally separated mentions of *pneumonia* were grouped together in the
2012 gold annotations but split into separate groups in the model output. Rather than a failure to track
2013 entities, this divergence suggests the model applies a stricter granularity, distinguishing findings based
2014 on specific attributes (e.g., diagnostic status or precise location) where human annotators might merge
2015 them. Similarly, regarding opacity in the right cardiophrenic sulcus, the model separated instances
2016 based on their evolving descriptions (e.g., “resolving”), prioritizing attribute precision over broad
2017 grouping.2018 Conversely, the model demonstrated the capability for semantic unification where appropriate. As
2019 seen in Patient p18517718, the model successfully reduced redundancy found in human annotations.
2020 For example, while the gold standard labeled specific attributes of a single medical device (e.g.,
2021 feeding tube, tip location) as separate groups, the model unified them into a single coherent entity. This
2022 highlights the model’s ability to identify core clinical concepts and organize fragmented descriptions
2023 effectively.2024 Overall, despite these variations in granularity, the grouping performance remained robust. The
2025 model preserved the essential semantic structure, balancing fine-grained distinctions for evolving
2026 pathologies with the integration of redundant observations. These findings support the reliability of
2027 our sequential annotation approach for tracking clinically meaningful entities over longitudinal report
2028 timelines.2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051

2052

2053

2054

2055

2056

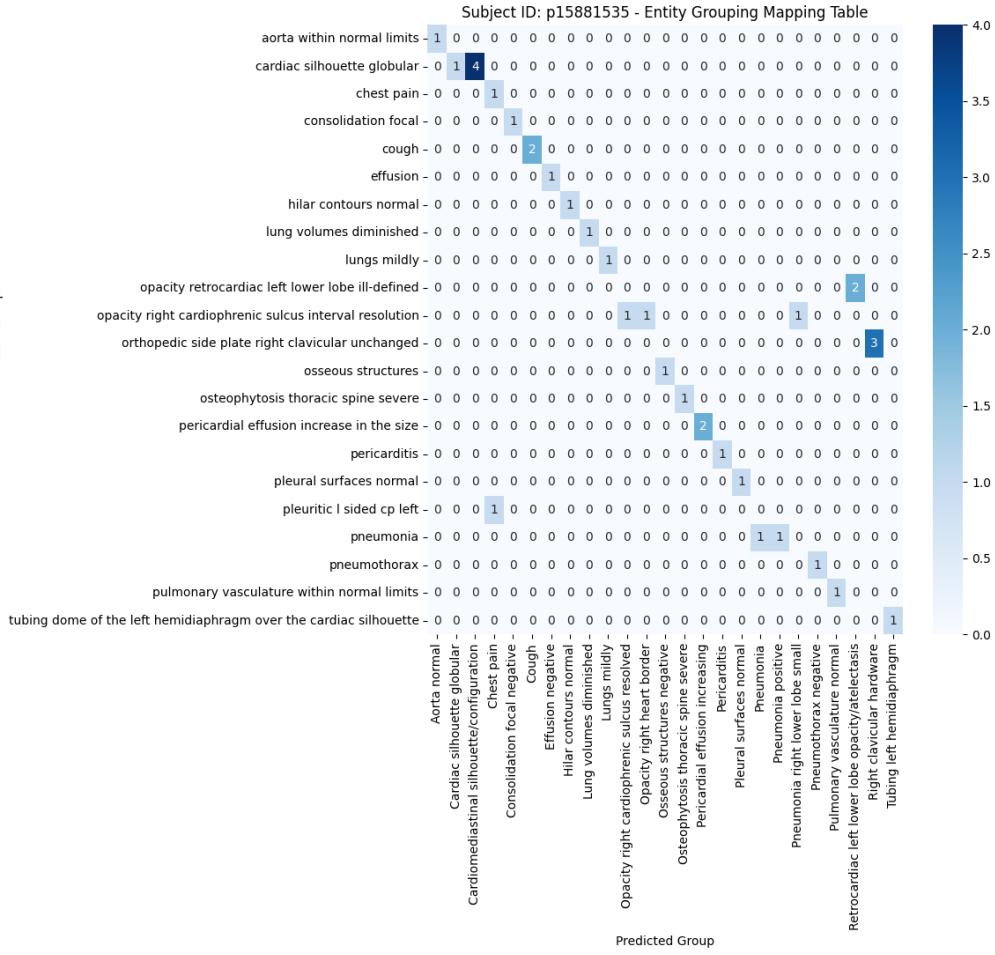
2057

2058

2059

2060

Actual Group



(a) Patient p15881535: Gold standard 22 groups vs Model 25 groups. Overall, the model aligns closely with the human gold standard. The slight difference in count stems from the model making more fine-grained distinctions—such as separating a diagnosis from its specific location (e.g., in Pneumonia or Cardiac Silhouette)—whereas human annotators tended to group them. These variations reflect a rigorous interpretation of the criteria while maintaining full semantic consistency.

Figure B.1: Entity grouping results for three sample patients based on sequential chest X-ray reports. The figures compare human-annotated gold-standard groupings (rows) with GPT-4.1 model predictions (columns). Numbered cells represent individual findings. Despite slight wording variations, the model demonstrates strong adherence to semantic grouping criteria. (Continued on next page)

2095

2100

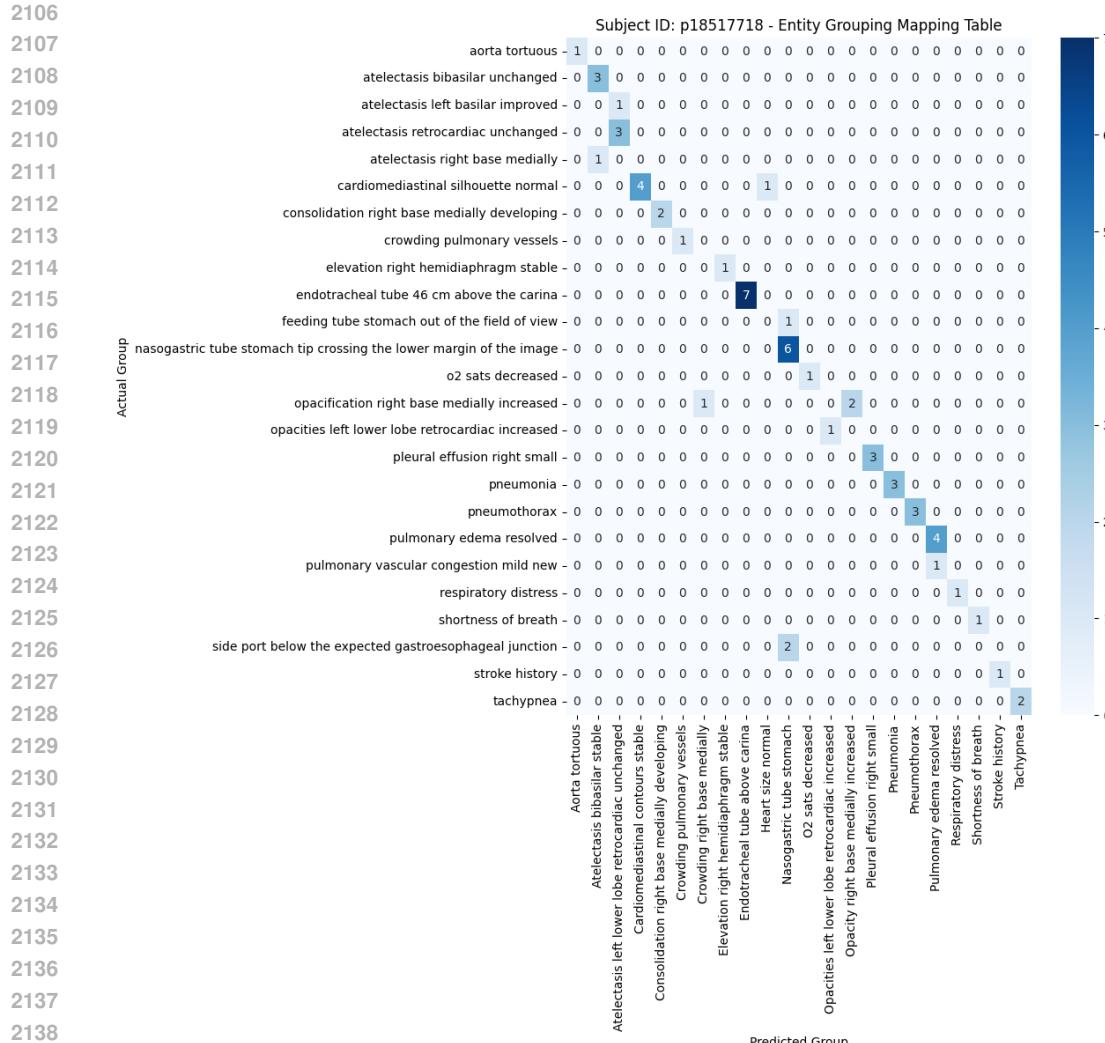
2101

2102

2103

2104

2105



(b) Patient p18517718: Gold standard 25 groups vs Model 22 groups. While the model achieves high overall alignment with the gold standard, subtle discrepancies arise due to differences in granularity levels regarding semantically overlapping entities. In contrast to the previous case, where human annotators separated specific attributes of a single medical device (e.g., feeding tube, tip location, side port), the model interpreted them as components of a single coherent entity ('Nasogastric tube stomach'). This illustrates that mismatches often result from legitimate variations in how granularly overlapping concepts are defined, rather than semantic errors.

Figure B.1: Continued.

2160 **C LUNGUAGESCORE DETAILS**
21612162 **C.1 ATTRIBUTE WEIGHTS OF LUNGUAGESCORE**
21632164 To reflect the clinical importance of structured attributes in radiology reports, LUNGUAGESCORE
2165 applies attribute-specific weights when measuring similarity between predicted and reference struc-
2166 tures. Each comparison is performed at the level of relational triplets, jointly assessing both temporal
2167 and structural alignment. For structural attributes, we assign weights based on expert consensus from
2168 the four board-certified radiologists who participated in the data annotation process, reflecting each
2169 attribute’s diagnostic significance. Although the initial weights are unnormalized, they are rescaled
2170 such that their total contribution sums to 1.0 during evaluation (see Table C.1).
21712172 For the sequential setting, temporal alignment contributes a fixed weight of 1.0, divided equally
2173 between two components: whether the predicted and reference findings belong to the same study
2174 timepoint (0.5), and whether they fall within the same temporal group (0.5).
21752176 Although our schema includes inferential relations such as `ASSOCIATE` and `EVIDENCE`, these are
2177 intentionally excluded from the evaluation metric. Such relations capture diagnostic reasoning—e.g.,
2178 linking “opacity” as supporting evidence for “pneumonia”—but do not directly reflect the correctness
2179 of factual information. Scoring them would conflate interpretive inference with structural accuracy.
2180 Instead, our metric focuses on clinically grounded descriptors and attributes that define the diagnostic
2181 content of the report. Future extensions may consider integrating reasoning-based relations in settings
2182 that explicitly target causal or explanatory fidelity.
21832184 Table C.1: Weights used in LUNGUAGESCORE for evaluating structural similarity. Temporal
2185 weights apply only in the sequential setting, while structural attribute weights reflect the diagnostic
2186 importance of each relation type. All values are normalized such that their respective groups (temporal
2187 or structural) sum to 1.0 during evaluation.
2188

		Structural Attribute Weights	Value
	DXSTATUS		0.50
	DXCERTAINTY		0.10
	LOCATION		0.20
	SEVERITY		0.15
	ONSET		0.15
	IMPROVED		0.15
	WORSENED		0.15
	PLACEMENT		0.15
	NO CHANGE		0.10
	MORPHOLOGY		0.05
	DISTRIBUTION		0.05
	MEASUREMENT		0.05
	COMPARISON		0.03
	PAST HX		0.01
	OTHER SOURCE		0.01
	ASSESSMENT LIMITATIONS		0.01

Temporal Weights	Value
Study Timepoint	0.5
Temporal Group	0.5

2204 **C.2 LUNGUAGESCORE EXAMPLES**
22052206 **Single-Report Assessment** To illustrate how LUNGUAGESCORE evaluates structured prediction
2207 quality in the single-report setting, we present detailed examples of pairwise comparisons between
2208 predicted and gold-standard structured reports. As detailed in Section 5 in the main text, each
2209 comparison is decomposed into two complementary components:
22102211

- **Semantic Score:** Computed as the cosine similarity between embedded linearized entity
2212 phrases. These phrases are formed by concatenating free-text attributes, including `LOCATION`,
2213 `MORPHOLOGY`, `DISTRIBUTION`, `MEASUREMENT`, `SEVERITY`, `ONSET`, `IMPROVED`,
2214 `WORSENED`, `NO CHANGE`, and `PLACEMENT`. This representation captures the semantic

2214 content of the entity and its descriptive qualifiers, allowing similarity to be measured in an
 2215 integrated manner.

2216

- 2217 • **Structural Score:** A weighted sum of attribute-wise comparisons. Categorical attributes
 2218 (DXSTATUS and DxCERTAINTY) are scored in binary fashion (1.0 for exact match, 0.0
 2219 otherwise), while all other attributes are evaluated via cosine similarity of their embeddings.
 2220 The relative importance of each attribute is determined by expert-defined weights (see
 2221 Table C.1).

2222 The final similarity between a predicted and reference finding is calculated as the product of the
 2223 semantic and structural scores:

$$2224 \text{TOTAL SCORE} = \text{Semantic Score} \times \text{Structural Score}$$

2225

2226 **Note:** Entity refers to the linearized phrase comprising the core entity and its attributes. Avg. Cosine
 2227 indicates cosine similarity averaged over MedCPTJin et al. (2023) and BioLORD23Remy et al.
 2228 (2024) embeddings of the phrases. Weights shown in the table reflect unnormalized values; the
 2229 final STRUCTURAL SCORE is computed by normalizing the weighted sum by the total weight of all
 2230 included attributes. For a more formal explanation of the scoring method, we refer to Section 5 in the
 2231 main text.

2232 Example 1: Moderate Match with Attribute-Level Divergence

2234 Attribute	2235 GT Value	2236 Pred Value	2237 Match Type	2238 Score	2239 Weight
Entity	effusions bilateral small	pleural effusion left-sided pleural small stable	Avg. Cosine	0.743	—
DxStatus	positive	positive	Exact match	1.00	0.50
DxCertainty	definitive	definitive	Exact match	1.00	0.10
Location	bilateral	left-sided pleural	Avg. Cosine	0.54	0.20
Severity	small	small	Exact match	1.00	0.15
Improved	—	stable	Avg. Cosine	0.00	0.15

2241 Semantic Score = 0.743, Structural Score = 0.681, Total Score = **0.506**

2242 Example 2: Partial Match with Location and Severity Differences

2246 Attribute	2247 GT Value	2248 Pred Value	2249 Match Type	2250 Score	2251 Weight
Entity	opacification left retrocardiac	pleural effusion left moderate	Avg. Cosine	0.447	—
DxStatus	positive	positive	Exact match	1.00	0.50
DxCertainty	definitive	definitive	Exact match	1.00	0.10
Location	left retrocardiac	left	Avg. Cosine	0.60	0.20
Severity	—	moderate	Avg. Cosine	0.00	0.15

2252 Semantic Score = 0.447, Structural Score = 0.758, Total Score = **0.339**

2253 Example 3: Strong Match with Minor Lexical Variants

2257 Attribute	2258 GT Value	2259 Pred Value	2260 Match Type	2261 Score	2262 Weight
Entity	opacity right lung base	opacity right lower lung base stable	Avg. Cosine	0.842	—
DxStatus	positive	positive	Exact match	1.00	0.50
DxCertainty	definitive	definitive	Exact match	1.00	0.10
Location	right lung base	right lower lung base	Avg. Cosine	0.95	0.20
Improved	—	stable	Avg. Cosine	0.00	0.15

2263 Semantic Score = 0.842, Structural Score = 0.902, Total Score = **0.759**

2264

2265 **Sequential-Report Assessment** To clarify how LUNGUAGESCORE computes similarity in the
 2266 sequential setting, we present illustrative examples comparing gold-standard and predicted findings.
 2267 Each score is computed from three components:

- **Semantic Score:** In the sequential-report setting, semantic similarity is computed between *ENTITYGROUP* representations, which group together lexically variable but conceptually equivalent findings observed at different timepoints.
- **Temporal Score:** Value of 1.0 if both findings appear in the same study timepoint and in the same TEMPORAL GROUP, or 0.5 if they belong to the same broader TEMPORAL GROUP but from different studies, or vice versa. If neither matches, the score is 0.
- **Structural Score:** Weighted average of attribute-level matches (exact for binary attributes, cosine similarity for textual ones).

The overall similarity score is computed as:

$$\text{Total Score} = \text{Semantic Score} \times \text{Temporal Score} \times \text{Structural Score}$$

Table C.2: Examples of LUNGUAGESCORE computations in the sequential setting. Each row compares a predicted finding against the corresponding ground-truth reference. Total Score is computed as the product of semantic similarity, temporal alignment, and structural accuracy. **Time** denotes the study timepoint, and **TG** indicates the assigned temporal group.

Case	GT			Prediction			Explanation	Total (Sem \times Temp \times Str)
	EntityGroup	Time	TG	EntityGroup	Time	TG		
1	pleural effusion subpulmonic moderate	2	1	pleural effusion right subpulmonic layering moderate stable	2	1	Minor semantic variation in anatomical modifiers and progression terms	0.68 (0.82 \times 1.0 \times 0.83)
2	hilar contours stable	3	1	hilar contours unchanged	3	1	Semantically equivalent; lexical variation in stability descriptor	0.90 (0.93 \times 1.0 \times 0.97)
3	atelectasis left lower lobe mild-to-moderate	1	1	atelectasis left lower lobe unchanged	2	1	Different timepoints (0.5), severity term vs. stability term mismatch	0.35 (0.92 \times 0.50 \times 0.76)
4	PICC mid SVC	2	1	left PICC mid SVC	1	1	Core entity match with modifier discrepancy; higher specificity in prediction; different timepoints	0.45 (0.90 \times 0.50 \times 1.00)
5	hilar contours unchanged	2	1	cardiomedastinal silhouette unchanged	3	1	Semantically related anatomical terms; timepoint mismatch (0.5)	0.34 (0.68 \times 0.50 \times 1.00)

Final Scoring and Interpretability LUNGUAGESCORE calculates a TOTAL SCORE for each matched pair of predicted and reference findings by combining semantic similarity and structural alignment. In the single-report setting, the total score is defined as the product of cosine similarity over linearized entity phrases and a weighted score of attribute-level matches. In the sequential setting, the metric further incorporates a temporal alignment factor, distinguishing between exact study-time matches and broader temporal group continuity.

These component-wise scores are then aggregated across matched pairs to compute the overall F1 metric, as detailed in Section 5. Crucially, each comparison yields interpretable diagnostics: the semantic score quantifies lexical alignment of free-text descriptors; the structural score exposes attribute-level agreement or divergence; and in longitudinal contexts, the temporal score reveals whether grouping decisions respect continuity over time.

By exposing this granularity, LUNGUAGESCORE not only delivers a robust scalar evaluation, but also supports nuanced error analysis—highlighting which components of a model’s output (e.g., misassigned severity, incorrect timing, lexical drift) most strongly influenced final performance. This interpretability makes the metric especially valuable to understand model’s behavior.

2322
2323

C.3 CLINICAL BERT MODEL SELECTION

2324
2325
2326
2327
2328
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338

We considered multiple clinical BERT models for computing contextual semantic embeddings. The candidate models we compared were BioLORD (Remy et al. (2024)), BiomedBERT (Gu et al. (2020)), MedCPT (Jin et al. (2023)), BioClinicalBERT (Alsentzer et al. (2019)), ClinicalBERT (Liu et al. (2025)) and BioBERT (Lee et al. (2020)). To decide which models to use in the semantic similarity step of LUNGUAGESCORE, we conducted an experiment over ReXVal, a subset of the MIMIC-CXR test set encompassing 50 randomly selected studies. We structured each individual study according to our framework described in Section 4(i), and then generated all linearized phrases derived from entity–location–attribute triplets for both the reference report and the candidate report. We then used each candidate BERT embedding model to generate an embedding for each phrase, and computed the pairwise cosine similarity for all pairs of phrases (one from the reference report and one from the candidate report). Figure C.1 shows the distribution of this similarity score for the different BERT embedding models. We find that BiomedBERT, BioClinicalBERT, ClinicalBERT and BioBERT lack variety, always scoring pairs of phrases as highly related. BioLORD manages to capture the most diversity in semantic similarity, followed by MedCPT. For this reason, we choose to use both BioLORD and MedCPT to calculate semantic similarity, by taking the average over both models.

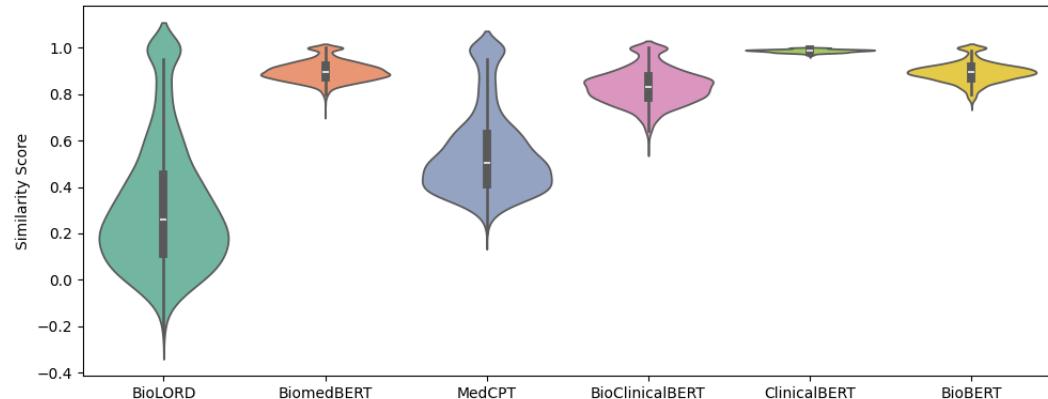
2339
2340
2341
2342
2343
2344
2345
2346
2347
2348
2349
2350
2351
23522353
2354
2355
2356
2357
2358
2359
2360
2361
2362
2363
2364
2365
2366
2367
2368
2369
2370
2371
2372
2373
2374
2375

Figure C.1: Distribution of pairwise cosine similarity scores for different BERT embedding models, calculated between pairs of embedded linearized phrases taken from the ReXVal dataset.

2376 D METRIC VALIDATION

2378 **Metric Implementation Details** Whenever not further specified, we used default settings for
 2379 all the metrics as provided by their respective libraries. For BLEU, we use the implementation
 2380 provided in the `huggingface/evaluate` library. For BERTScore, we also use the implementation
 2381 from the `huggingface/evaluate` library, with `distilroberta-base` as an embedding model. For
 2382 GREEN, we use `StanfordAIMI/GREEN-rad11ama2-7b` as a language model. For FineRadScore,
 2383 we use GPT-4 as a language model, which responds with a list of errors each linked to a severity
 2384 level. To turn this into a score, we associate each severity level with a number, and sum these scores,
 2385 forming FineRadScore as proposed by (Huang et al. (2024)). In our tables, we report 1/FineRadScore,
 2386 inverting the total sum to ensure that a higher score is associated with higher quality. For RaTEScore,
 2387 we use their default weight matrix. Note that in their own comparison with ReXVal, the authors used
 2388 a custom weight matrix trained specifically for long reports instead of the default, explaining the
 2389 slight discrepancy between their reported Kendall Tau correlation with ReXVal radiologists and the
 2390 one we report in Table 2. We also report results for RadGraph and its newer version RadGraph-XL,
 2391 using the RG_EG setting to calculate the F1 score as proposed in (Delbrouck et al. (2022)).

2392 **ReXVal Analysis** To assess the consistency of our metric with established evaluation standards,
 2393 we conducted a correlation analysis across the ReXVal benchmark, which includes expert-annotated
 2394 radiology reports and associated error counts. Specifically, we computed pairwise Pearson cor-
 2395 relations between all single-report metrics over the ReXVal dataset. As presented in Figure D.1,
 2396 our metric exhibits strong positive correlations with BLEU (0.73), BERTScore (0.77), GREEN
 2397 (0.84), RaTEScore (0.77), 1/FineRadScore (0.73), RadGraph F1 (0.80) and RadGraph-XL F1 (0.80).
 2398 Notably, among all evaluated metrics, our score achieves the highest average correlation across
 2399 all pairwise comparisons, indicating strong alignment with multiple evaluation perspectives and
 2400 suggesting broader generalizability.

2401 Furthermore, Figure D.2 illustrates the linear relationship between each metric and the number of
 2402 radiologist-identified errors per ReXVal report. Although 1/FineRadScore shows the highest overall
 2403 correlation, its relationship with error counts is not consistently linear, especially when the number of
 2404 errors is low. In these cases where distinguishing between high-quality outputs is most crucial, its
 2405 ability to make fine-grained distinctions is limited. In contrast, our metric not only maintains strong
 2406 correlation but also demonstrates stable linear responsiveness across the full error range, underscoring
 2407 its robustness and reliability as a clinically aligned evaluation measure.

2408
 2409
 2410
 2411
 2412
 2413
 2414
 2415
 2416
 2417
 2418
 2419
 2420
 2421
 2422
 2423
 2424
 2425
 2426
 2427
 2428
 2429

2430

2431

2432

2433

2434

2435

2436

2437

2438

2439

2440

2441

2442

2443

2444

2445

2446

2447

2448

2449

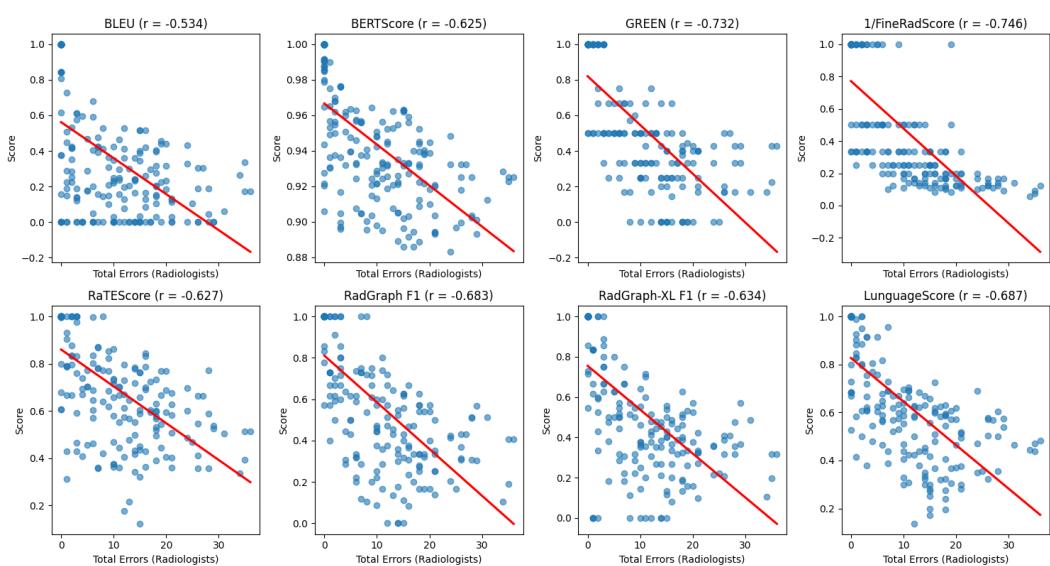
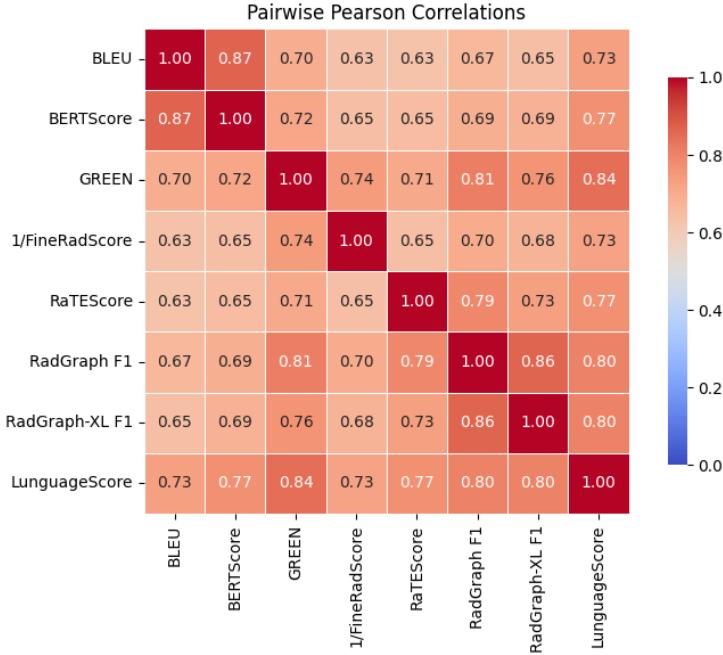
2450

2451

2452

2453

2454 Figure D.1: Pairwise Pearson correlations between our metric (LUNGUAGESCORE), and the metrics
 2455 BLEU, BERTScore, GREEN, 1/FineRadScore, RaTEScore, RadGraph F1 and RadGraph-XL F1.



2479 Figure D.2: Scatter plot illustrating the correlation between the total number of errors identified by
 2480 radiologists per report, and each of the single-report metrics, including our LUNGUAGESCORE. r
 2481 indicates the Pearson correlation as reported in Table 2.

2482

2483

2484
 2485 **Error Sensitivity Analysis with ReXErr (Rao et al. (2024))** To assess the error sensitivity of our
 2486 metric across diverse failure types in radiology report generation, we use the ReXErr-v1 dataset (Rao
 2487 et al. (2025)), which contains synthetic reports with systematically injected clinical errors. These
 2488 errors are categorized into content addition, context-dependent, and linguistic quality types, covering
 2489 a broad spectrum of realistic mistakes. We focus on the subset of ReXErr aligned with our sequential
 2490 structured report dataset, comprising 57 MIMIC-CXR reference reports paired with corresponding
 2491 error-injected versions. Each manipulated report contains three injected errors, drawn from 12 defined
 2492 error categories using a context-sensitive sampling method.

2493 For each pair, we extract the Findings and Impression sections and evaluate them independently using
 2494 our single-report LUNGSCORE, along with established alternatives: GREEN, FineRadScore,
 2495 and RaTEScore. Figure D.3 displays the score distributions for each of the 12 error types, relative to
 2496 the average score across the subset. Our metric demonstrates differentiated sensitivity across error
 2497 types, with notably larger penalizations for false predictions, incorrect negations, and changes in
 2498 severity—reflecting its alignment with clinically meaningful deviations.

2499 **Sequential Sensitivity Analysis** We further assessed the sensitivity of LUNGSCORE to
 2500 clinically meaningful disruptions in temporal coherence by constructing a synthetic evaluation set in
 2501 which longitudinal progression cues were deliberately inverted. Specifically, we selected 8 patient
 2502 sequences from our sequential-report dataset that contained explicit temporal descriptors—such as
 2503 *improved* or *worsened*—and manually reversed these attributes to simulate a contradiction in the
 2504 clinical trajectory. For example, a statement like “the previously seen right lower lobe opacification
 2505 has decreased substantially” was changed to “increased substantially,” thereby inverting its semantic
 2506 implication. Two patient sequences that lacked any such temporal expressions were excluded.

2507 Both the single-report and sequential variants of LUNGSCORE were applied to these perturbed
 2508 sequences. To quantify the metric’s responsiveness, we introduce the *Effect Rate*, which captures the
 2509 average score reduction per flipped attribute:

$$2510 \text{Effect Rate (\%)} = \frac{1 - \text{score}}{\#\text{flipped attributes}} \times 100$$

$$2511$$

$$2512$$

2513 A perfect score of 1.0 indicates complete semantic and structural agreement with the gold standard.
 2514 Deviations from this ideal reflect the metric’s sensitivity to reversed temporal directionality. The
 2515 normalization by the number of flipped attributes allows us to measure the per-attribute impact on the
 2516 similarity score.

2517
 2518 Table D.1: Effect Rate for each manipulated patient sequence. W/I denotes the number of
 2519 worsened/improved attributes flipped.

Patient ID	# Attr. (W/I)	Single Score	Effect Rate (S, %)	Sequential Score	Effect Rate (Seq, %)
p10274145	5 (0/5)	0.981	0.38	0.979	0.42
p10523725	3 (1/2)	0.989	0.37	0.987	0.43
p10886362	8 (5/3)	0.983	0.21	0.979	0.26
p10959054	13 (9/4)	0.967	0.25	0.963	0.28
p12433421	15 (8/7)	0.968	0.21	0.971	0.19
p15321868	2 (1/1)	0.982	0.90	0.988	0.60
p15881535	1 (0/1)	0.992	0.80	0.992	0.80
p18079481	10 (2/8)	0.976	0.24	0.980	0.20

2520 While the absolute Effect Rates are relatively small (typically below 0.5%), they scale proportionally
 2521 with the number of flipped attributes, indicating that LUNGSCORE reliably captures the semantic
 2522 impact of trend reversals. Notably, even sequences with a single flipped term exhibited pronounced
 2523 per-attribute degradation, highlighting the metric’s granularity and responsiveness. These results
 2524 affirm that LUNGSCORE can effectively detect inconsistencies in longitudinal directionality,
 2525 even when the surface fluency of the report remains intact.

2526
 2527
 2528
 2529
 2530
 2531
 2532
 2533
 2534
 2535
 2536
 2537

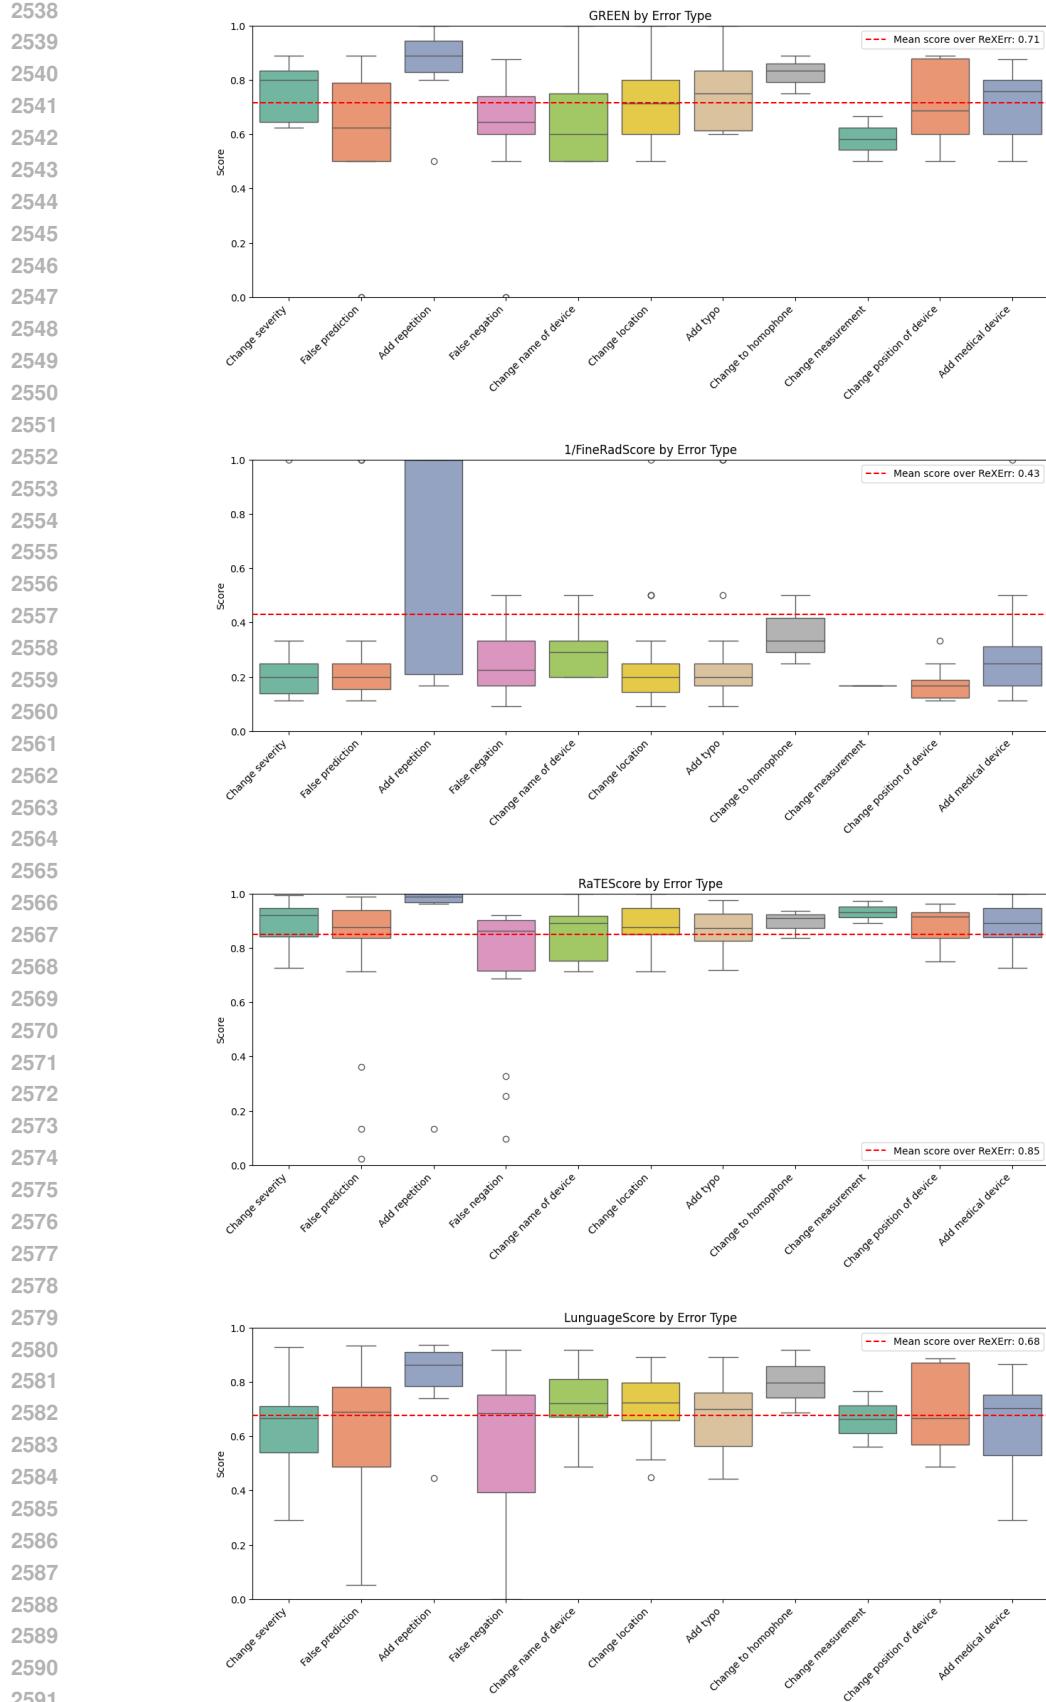


Figure D.3: Distribution of the scores for each of the twelve error types in ReXErr, relative to the average score across the 57 ReXErr reports.

2592 **E SYNTHETIC REPORT GENERATION DETAILS**
 2593

2594 **MAIRA-2 (Bannur et al. (2024))** At the input, we feed in a frontal chest X-ray image for the
 2595 current study. If there is no frontal available for the patient, we do not generate a report. If there are
 2596 multiple frontals, we randomly choose one. We also pass along a random lateral chest X-ray image
 2597 for the current study, should it be available. MAIRA-2 additionally accepts the indication, technique
 2598 and comparison sections. We therefore input the history for the current study in the “indication” field,
 2599 if there is one. For the comparison, we input “Chest radiography dated _.” if there is a previous study,
 2600 to comply with the anonymised dates in the MIMIC-CXR dataset. We do not input a technique, since
 2601 this field could not be reliably extracted for the MIMIC-CXR test set. We explore two distinct ways
 2602 for including prior information in the generation setup. In the **standard** setting, we input the ground
 2603 truth reference report that is available for the previous study. This report is structured following
 2604 the template “INDICATION: <prior_history> COMPARISON: <prior_comparison> FINDINGS:
 2605 <prior_findings> IMPRESSION: <prior_impression>.”, where <prior_history>, <prior_impression>
 2606 and <prior_findings> are all taken from the previous study’s ground truth reference report, and
 2607 substituted by “N/A” if they are missing. If there is no prior study, the prior report field is set to
 2608 “None”. If the previous study was the first one in the sequence, then <prior_comparison> is set to
 2609 “N/A”, otherwise it is set to “Chest radiograph dated _.” In the **cascaded** setting, <prior_findings>
 2610 is set to the findings report that was generated in the previous study (if there is one, otherwise the
 2611 prior report field is set to “None”), while <prior_impression> is left blank (because MAIRA-2 only
 2612 generates findings), and the other inputs remain the same. In both settings, we input the frontal view
 2613 from the prior study, if there is one, and if there are multiple options, we choose the same one that
 2614 was used to generate the previous report. We ask MAIRA-2 to generate the findings section for the
 2615 current study, using their default settings, without grounding.

2616 **Medversa (Zhou et al. (2024))** Next to the current frontal image, we also fill in the additional input
 2617 fields expected by Medversa, which are context, prompt, modality and task. For context, we follow
 2618 the template “Age: None. Gender: None. Indication: <current_history>”. For <current_history>,
 2619 we pass along the “history” section of the reference report, should it be available, and otherwise we
 2620 set it to “None”. The modality and task are set to “cxr” and “report generation” respectively. All
 2621 language generation parameters are left as default. The prompt is set to “Can you provide a report of
 2622 <img0> with findings and impression?”. Note that this is the only model with the ability to generate
 2623 an impression section, and it will therefore naturally have an advantage over the other models when
 2624 we compare it to the reference report, where both the findings and impression section are included
 2625 based on their availability in the ground truth.

2626 **LIBRA (Zhang et al. (2025))** For LIBRA, a general-purpose medical vision–language model, we
 2627 use the authors’ public implementation in its image–text report generation mode. For each study,
 2628 we provide the current frontal chest X-ray as the primary input image; if multiple frontal views are
 2629 available, we randomly select one, and if no frontal is available, we do not generate a report. When
 2630 a previous study exists, we additionally pass the frontal image from the most recent prior study as
 2631 a second input so that LIBRA can jointly attend to the current and prior examinations; otherwise,
 2632 only the current image is used. We use a fixed, generic instruction prompt asking the model to
 2633 produce a detailed description of the radiographic findings, and keep all decoding hyperparameters at
 2634 their default values. The resulting text is taken as the model’s report for evaluation without further
 2635 post-processing or templating.

2636 **RGRG (Tanida et al. (2023)), Cvt2DistilGPT2 (Nicolson et al. (2023)), MedGemma (Sellergren
 2637 et al. (2025)), ChexAgentChen et al. (2024), and Lingshu (Xu et al. (2025))** For these models,
 2638 we input the current frontal chest X-ray image, randomly selecting one when multiple views were
 2639 available and skipping generation when none were present. We used the vision–language variant
 2640 google/medgemma-27b-it for MedGemma and the MIMIC-CXR-trained version with default
 2641 configuration for Cvt2DistilGPT2. RGRG and Cvt2DistilGPT2 generated only *findings* without a
 2642 separate *impression* section, whereas MedGemma and Lingshu produced full reports containing both
 2643 *findings* and *impressions*. The exact prompt templates for MedGemma and Lingshu, as specified in
 2644 their original papers, are shown below.

2646

2647

2648

2649

2650

2651

2652

2653

2654

2655

2656

2657

2658

Prompt Templates for MedGemma and Lingshu

[MedGemma]

You are an expert radiologist. Please succinctly describe the findings for the above chest X-ray.

[Lingshu]

You are a helpful assistant. Please generate a report for the given images, including both findings and impressions.

Return the report in the following format:

Findings: {} Impression: {}

F LIMITATIONS AND FUTURE DIRECTIONS

While LUNGUAGE defines the fine-grained evaluation dataset, the structuring framework produces schema-aligned representations, and LUNGUAGESCORE provides the scoring function for both single-report and longitudinal evaluation, the current work also has several limitations that suggest concrete directions for extension.

First, although LUNGUAGE provides fine-grained entity-, attribute-, and longitudinally interpreted annotations, its patient coverage remains modest: the longitudinal subset currently comprises 30 patients and 186 reports drawn from a single public MIMIC-CXR dataset. Both the dataset and the underlying schema were developed on this subset, so the current vocabulary and relation set may underrepresent findings, reporting conventions, and temporal patterns present in other institutions or populations. To support broader use, future work should scale and stress-test the schema and pipeline on the full MIMIC-CXR dataset as well as other large longitudinal datasets, and develop complementary benchmarks on multi-center, multi-country cohorts and additional imaging modalities. Because the schema, prompts, and implementation are publicly released, researchers can adapt the framework to local reporting conventions, extend the taxonomy, or substitute alternative structuring models while retaining a comparable evaluation protocol. The same tooling can also be used to generate large “silver-standard” structured sets from unlabeled reports, enabling end-to-end pipelines where models are trained and evaluated under a consistent schema, and to derive downstream resources such as QA or instruction-tuning datasets grounded in structured longitudinal trajectories. We hope this work will serve as a starting point for a broader community effort toward clinically grounded, temporally aware evaluation standards for radiology report generation.

As a second limitation and direction for future work, we note that advancing patient-centered reporting will require integrating structured EHR information alongside chest X-rays, including laboratory data, vital signs, procedures, and free-text clinical notes. Current image-based generation approaches struggle with context-rich sections such as patient history, and models that lack access to these contextual signals remain fundamentally limited in longitudinal reasoning and diagnostic continuity. A natural next step is therefore to extend our schema, structuring framework, and LUNGUAGESCORE to multimodal trajectories that couple images, reports, and EHR data, and to examine how the same design principles can be adapted to other imaging domains and healthcare settings. Key challenges for this extension include defining clinically reliable ground truth for multimodal trajectories, aligning heterogeneous temporal signals across modalities, and ensuring that extended versions of LUNGUAGESCORE remain interpretable and robust at EHR scale.

Third, in this work our structuring and evaluation operate purely at the report level. Although the schema explicitly distinguishes image-groundable entities (for example, perceptual findings and devices) from non-chest x-ray findings (for example, clinical history or laboratory results), we do not yet link these entities to the underlying images. An important next step is to ground LUNGUAGE in the pixel space by associating structured findings with spatial annotations, such as view-specific bounding boxes or pixel-level masks for lesions, devices, and other relevant regions. This would enable joint evaluation of whether a generated report is not only semantically and temporally consistent, but also spatially aligned with the visual evidence, and would support the construction of downstream vision-language tasks such as grounding, question answering, and instruction tuning based on the same structured representation.

2698

2699