

# RadTimeline: Timeline Summarization for Longitudinal Radiological Lung Findings

Sitong Zhou, Meliha Yetisgen, Mari Ostendorf

University of Washington  
{sitongz, melihay, ostendorf}@uw.edu

## Abstract

Tracking findings in longitudinal radiology reports is crucial for accurately identifying disease progression, and the time-consuming process would benefit from automatic summarization. This work introduces a structured summarization task, where we frame longitudinal report summarization as a timeline generation task, with dated findings organized in columns and temporally related findings grouped in rows. This structured summarization format enables straightforward comparison of findings across time and facilitates fact-checking against the associated reports. The timeline is generated using a 3-step LLM process of extracting findings, generating group names, and using the names to group the findings. To evaluate such systems, we create RadTimeline, a timeline dataset focused on tracking lung-related radiologic findings in chest-related imaging reports. Experiments on RadTimeline show tradeoffs of different-sized LLMs and prompting strategies. Our results highlight that group name generation as an intermediate step is critical for effective finding grouping. The best configuration has some irrelevant findings but very good recall, and grouping performance is comparable to human annotators.

**Keywords:** structured summarization, longitudinal EHRs, LLMs

## 1. Introduction

Tracking changes in status of findings in longitudinal radiology reports is critical for understanding disease progression, treatment response, and supporting disease risk assessment. Reading longitudinal reports can capture subtle changes that may be missed in isolated reports. Individual reports often document changes in status with short phrases (e.g., "improved", "persists") (Khanna et al., 2023; Park et al., 2024), but the longitudinal sequence of reports can provide details that make this information more useful. However, reading longitudinal records and pulling together related findings can be time consuming and contributes to cognitive overload. Artificial intelligence (AI) can support this work, both in the identification and presentation of relevant findings.

While AI tools for summarization are effective for many tasks, most work on multiple documents has involved unstructured summaries, which complicates fact-checking of statements about temporal changes (Chien et al., 2024). For high-stakes clinical scenarios, the ability to confirm source details behind an AI statement is critical. Verma et al. (2025) propose a verifiable structured summarization format in which each fact is linked to its source document and organized into human-curated topics. A pre-defined set of topics is useful for large document collections, but this is not practical for summarizing radiological findings for specific patients, where it may be useful to track multiple nodules observed in the same region of the lung, for example.

In this paper, we introduce a new **two-**

**dimensional timeline format for structured summarization** that makes it easy to see temporal changes. As shown in Figure 1 for lung-related findings, each column corresponds to a time-stamped radiology exam (as indicated in the header), such as a chest CT taken in May of a given year. Each cell in a column corresponds to a lung finding fact from that report, including details relevant to lung cancer risk assessment (e.g., size, location, status, etc.), such as "stable subsolid pulmonary nodule in the right upper lobe." Each row represents a group that tracks the evolution of a particular finding over time, with a row header describing the distinguishing characteristics of that group.

To support research with this framework, we create **RadTimeline, an evaluation dataset** with lung-related radiological findings within longitudinal chest-related imaging reports where each patient is associated with a hand-corrected timeline. In addition, metrics are proposed for assessing automatically generated timelines in terms of finding factuality, group name quality, and finding grouping performance.

We propose a **3-step LLM-based approach to timeline generation** that: (i) extracts lung-related finding facts from individual reports (the cells in a column), (ii) generates interpretable group names (row headers) from longitudinal reports, and (iii) associates each finding with its corresponding group (row).

In experiments with RadTimeline, we show that this 3-step timeline generation approach detects most lung findings correctly, but further work is needed to address irrelevant findings. We observe that group name quality is critical for effective group-

2 Group Name Generation Q: List all distinct group names given the longitudinal history.		1 Finding Extraction Q: What are lung findings from this report?	
group name	[YYYY-1]-04_chest_ct	[YYYY-1]-08_chest_ct	... YYYY-05_chest_ct
<i>right upper lobe ground glass nodule</i>	alveolar groundglass opacity in the posterior segment of the <b>right upper lobe</b> .. measuring 10.2 mm	<b>Right upper lobe</b> , indistinct mostly groundglass-like nodular lesion..., 10 mm overall size, <b>unchanged</b>	<b>stable</b> subsolid pulmonary nodule in the <b>right upper lobe</b>
<i>groundglass opacity in the posterior basal segment of the right lower lobe</i>	predominantly <b>groundglass opacity</b> in the posterior segment of the <b>right lower lobe</b> ... measuring as much as 2.0 cm	<b>Right lower lobe</b> , posterior basal segment, stellate-like groundglass opacity, 2 cm, <b>unchanged</b>	<b>stable</b> subsolid pulmonary nodule in the <b>right lower lobe</b>
<i>peripheral subsolid nodule in right upper lobe</i>			<b>new</b> 12 mm peripheral subsolid opacity in the <b>right upper lobe</b> anterior laterally, ...

**3 Group Assignment**  
Q: Which is the best group for this finding?

Figure 1: The timeline task and a three-step LLM approach. Each column corresponds to a time-stamped radiology exam (e.g., YYYY-05\_chest-ct). Each cell in a column is a piece of lung finding fact (e.g., "stable subsolid pulmonary nodule in the right upper lobe") from that report including clinically important details (e.g., "stable", "subsolid"). Each row groups temporally related findings, with a row header describing the distinguishing characteristics of that group (e.g., "right upper lobe ground glass nodule").

ing, particularly for weaker LLMs and embedding-based grouping strategies. The best prompting configuration reaches near-human performance in grouping findings.

## 2. Literature Review

### 2.1. Radiology report summarization

LLMs have been applied to summarize the findings sections of individual radiology reports into impressions, effectively capturing critical findings. Van Veen et al. (2023) apply LoRA fine-tuning to CLIN-T5-LARGE, and human evaluations show that the model generally captures critical findings but occasionally hallucinates details or infers prior history not in the report. Van Veen et al. (2024) demonstrate GPT-4 with in-context examples can generate more complete summaries with fewer errors than expert-written impressions, according to human judgment. Zeng et al. (2024) add in-context examples for Llama 3 70B, and observe improved automatic metrics such as ROUGE-L, but more factual errors according to human judgment. Revising these summaries with an LLM can correct presence errors but has minimal effect on errors related to progression status.

All the works mentioned above report automatic summarization metrics (BLEU, ROUGE-L, BERTScore) together with human evaluation, such as Likert-scale ratings on completeness and factuality (Van Veen et al., 2023). However, those automatic summarization metrics can correlate poorly with human judgments (Van Veen et al., 2023). For

LLM text generation evaluation, fine-grained factuality evaluation systems (Min et al., 2023; Yang et al., 2024) are proposed to measure on atomic fact level. They decompose generated text into atomic facts and score each fact individually.

For longitudinal report summarization, Chien et al. (2024) use LLMs to generate summaries directly from multiple reports. However, these unstructured summaries are difficult to verify, may omit findings and key details, and do not necessarily describe temporal changes for each finding. Structured summarization has been developed for heart failure (Verma et al., 2025), linking each fact in the summary to its source document and organizing facts into human-curated, disease-specific topics. However, fixed topic structures are not well suited for radiological findings, as groups of temporally related findings vary across patients.

Different from prior work, we develop a longitudinal report summarization approach to convert a patient's longitudinal chest imaging reports (CT, X-ray) into a timeline table, leveraging off-the-shelf LLMs and embedder models without using supervised labels. We first extract atomized lung-finding facts from the raw reports, then group temporally related facts to display the finding trajectories. Like some earlier work, we use ROUGE-L to evaluate findings, but we also provide scores for overgeneration of findings, quality of group names, and grouping of findings. Lastly, we compare automatic and human scoring for one LLM configuration.

## 2.2. Cross-document Coreference Resolution and Clustering

Our tasks share similarity with Cross-document coreference (CDCR) and clustering tasks, and their related methods that inspire our approach.

**CDCR** entails detection of coreferring event mentions from multiple documents. Our lung radiological findings can be considered as events, but we group temporally related finding facts, instead of coreferring mentions. CDCR is often evaluated with same-document coreference metrics, MUC, B<sup>3</sup>, CEAF, their average CoNLL F1 (Pradhan et al., 2012). Barhom et al. (2019) develop supervised neural models that model mention pairs. Caciularu et al. (2021) further improve the supervised results by leveraging cross-document contexts using Longformer-based models.

LLMs have been used directly for CDCR without supervised learning and require prompt engineering. Zhao et al. (2023) use zero-shot GPT-4 to classify pairs of decontextualized sentences of event triggers that contain document-level information, outperforming untrained crowd workers from MTurk. Min et al. (2024) use GPT-4 to assign cluster indexes to mentions and output in a JSON format, and find that GPT-4 is worse than their supervised method, and using full contexts from all documents is worse than using only sentences containing the mentions. Sundar et al. (2024) ask LLMs to replace placeholders next to the mentions with major entity tags in long narratives. When evaluated on gold mentions, GPT-4 is slightly better than their supervised method but not GPT3.5.

For radiology finding tracking, Datta et al. (2022) present a CDCR dataset about tracking findings and devices across each patient’s longitudinal radiology reports, and fine-tune a BERT model for pairwise mention coreference achieving low-to-moderate CoNLL F1 scores. Mathai et al. (2025) use zero-shot LLMs to match a sentence in a follow up report to the single best matching sentence in its prior report.

Our group assignment step is similar to the radiology CDCR task presented by Datta et al. (2022), but differs in that we do not group elements via pairwise associations. Instead, we create groups by mapping elements to a group label using LLMs, either via a question-answering prompt or using a prompt that inserts group tags for sequence of findings, inspired by Sundar et al. (2024). We also investigate the impact of full longitudinal contexts, given the mixed results in previous studies (Caciularu et al., 2021; Min et al., 2024).

Our task is also similar to **Clustering** in terms of grouping similar items, where items are findings and similarity is in terms of shared characteristics among temporally related findings. Off-the-shelf embedders can provide similarity measures for

clustering items. Muennighoff et al. (2023) present a MTEB benchmark to evaluate text embedders on varying tasks and observe that no single embedder dominates all tasks. E5-mistral-7b (Wang et al., 2024) is a finetuned Mistral-7B bi-encoder trained on diverse synthetic data and among the state-of-art embedder models on MTEB. They support prepending textual instructions to queries for task adaptation.

LLMs can guide clustering by infusing task-specific preferences. Huang and He (2024) use a two-stage LLM pipeline built on GPT-3.5. They first ask LLM to generate human-interpretable group names, followed by a de-duplication step, then classify each item into a group. They find the performance gap between using generated groups and using oracle group names is small. This is very relevant to our LLM approaches for the group name generation and group assignment steps. Our work differs in its focus on long findings for a single patient vs. corpus-level topics. We also collect group name reference data to evaluate the intermediate name generation step.

## 3. Gold Timeline Dataset Curation

We sample 10 patients from a patient cohort from an existing case-control dataset created for a lung cancer surveillance project in UW Medicine,<sup>1</sup> where temporal changes of lung-related findings are clinically important risk factors. We ensure that each of the sampled patients having at least four chest-related radiology exams, including chest CT, chest X-ray, and abdomen CT, within a five-year window. According to the state cancer registry, 6 patients have future lung cancer diagnosis records within 3 years, and 4 patients have no diagnosis records of any type of cancer. The raw EHRs contain 65 chest-related exam reports from all patients. The number of reports per patient ranges from 4 to 12, and longitudinal report lengths measured in tokens range from 646 to 4625, with an average at 2055.

We create a timeline for each patient based on their longitudinal radiology reports, as illustrated in Figure 1. Each column represents the findings (**gold findings**) extracted from an individual radiology report, while each row captures the temporal evolution of a specific finding mentioned across multiple reports (**gold group assignments**). Each row is further assigned a **gold group name**, summarizing the key characteristics of the findings in this row. Since our focus is on tracking lung-related conditions, findings pertaining to other organs (e.g., the aorta) or descriptions of medical devices are considered irrelevant and are omitted from the timelines. We work with two medical students and one

<sup>1</sup>A subset of this cohort was previously used by Zeng et al. (2025) for lung cancer prediction.

medical doctor as annotators during the timeline creation process. We first automatically generate timelines using LLMs (see Section 4 for details) and hand-correct LLM generated findings and group assignments based on the feedback from two medical student annotators. Finally, a third annotator reviews the gold finding groups, and hand-corrects row names of each version of gold group assignments to produce the gold group names.

Annotation guidelines are included in Appendix A.

### 3.1. Gold Findings (Columns)

**Factuality evaluation:** We ask two medical student annotators to rate the factuality of each generated finding against its associated report, and to add any missing findings that the model fails to include. In addition, they provide notes explaining the identified errors, which are later used by an author to manually correct the generated findings as the gold-standard reference.

To reduce annotation efforts and focus on abnormalities, we manually drop absent findings, such as "no new nodules," and normal findings, such as "clear lungs." The original automatically generated timelines contain 221 findings from 61 reports. After removing absent findings, we obtain 155 findings from 55 reports for the annotation tasks in factuality evaluation.

We define four categories for factuality rating: "irrelevant," "wrong," "partially correct," and "correct." "Irrelevant" refers to findings that are not related to lung radiological observations, typically describing findings from other organs or medical devices. "Wrong" refers to findings that contain errors regarding presence or temporal change trends. "Partial" is used for findings that contain inaccurate or incomplete details. For example, in "stable sub-solid pulmonary nodule in the right upper lobe", annotators will consider it as "partial" if "stable" or "subsolid" are missing. Findings that include correct and complete clinically important details are labeled as "Correct." Table 1 shows the confusion matrix between two annotators' ratings. The two annotations reach a high agreement at 86%. No finding is considered wrong by both of the annotators, and the major mistake is including irrelevant findings (14%). In addition, one annotator identifies 3 missing findings from a single report. No other missing findings are reported.

**Gold Finding Curation:** We refine the LLM-generated findings as follows. Irrelevant or wrong findings are removed, correct findings are kept unchanged, and partially correct findings are either revised by adding missing details, or merged with another finding with the complementary details when they are both described in proximity within the same

	Wrong	Partial	Correct	Irrelevant
Wrong	0	0	2	0
Partial	0	12	3	0
Correct	2	9	106	2
Irrelevant	1	0	2	16

Table 1: Confusion-matrix of factuality annotation between two annotators. Annotators agree on factuality ratings for 86% of 155 generated findings.

paragraph. Additionally, we include the three missing findings identified during review. If annotators rate a finding differently, the first author resolves the disagreements based on annotators' notes to decide the action. One exception is made for one finding labeled as correct by both annotators but later found to be irrelevant during post-analysis and removed from the gold references. For four partially correct findings in which incomplete details result from complementary information spread across different parts of a radiology report, we create a new finding and add a same-report link between the new finding and the original partial finding during the subsequent group assignment curation process. This is consistent with the automatically generated findings, where finding details from different parts of reports are not always merged.

After resolving disagreements, we find that 15% of the 155 findings used in the factuality human evaluation can be dropped due to being wrong or irrelevant and 12% could be further merged with another finding in the same report. After removing and adding findings based on ratings and notes from annotators, there are 136 gold findings, or 115-116 merged findings after merging same-report linked findings grouped separately by the two annotators.

### 3.2. Gold Group Assignments (Rows)

To collect group annotations for temporally related findings, annotators assign a group ID to each finding. We provide them with automatically generated group names linked to the corresponding group IDs, displayed as the row headers in the generated timelines. Annotators may add new group names and IDs if the existing list is incomplete or not representative of a group. When two group IDs convey the same meaning, annotators are instructed to select one ID and apply it consistently across all relevant findings. Revising the group names is not required at this stage. Annotators have access to the full longitudinal radiology reports while performing group ID assignment. For the seven added findings (3 missing, 4 due to missing details), group IDs have not been pre-assigned. To assign IDs to these findings, one annotator labels the seven findings in their own group assignment version, and

Statistic	Ann-1	Ann-2
<b># Groups</b>	65	62
Avg. / Max size	2.1 / 8	2.2 / 7
% Non-singleton	52	50
<b># Findings (Total # 136 )</b>		
% Linked (Any / Cross)	77 / 70	77 / 71
<b># Findings ( Merg.)</b>	116	115

Table 2: Group annotation statistics from two annotators (Ann-1 and Ann-2) on findings in RadTimeline. The group size is based on the 136 findings. "Findings (Merg.)" are findings that are merged with same-report links for finding extraction evaluation. The % Linked refer to the percentage of findings that have links to any other finding (Any) vs. only cross-note findings (Cross) among the total 136 findings.

the first author propagates these labels to another annotator’s version to maintain consistency.

We collect two versions of group assignments to capture individual annotator preferences. For example, for "low lung volumes bilaterally", and "bibasilar opacities, likely atelectasis", one annotator groups them and the other separates them into different groups. The CoNLL F1 between two gold group assignments is 82, reflecting good inter-annotator agreement. Group annotation statistics are in Table 2.

### 3.3. Gold Group Names (Row Headers)

The third annotator who is a medical doctor reviews both versions of gold finding groups. The annotator is asked to manually correct row names of each version of gold group assignments to produce distinguishable group names, and each name should cover the shared information within the group.

## 4. Methods

We use a three-step approach to generate timelines. First, lung-related findings are identified for each report. Second, given the longitudinal report contexts, a list of group names is generated to represent distinct findings. Lastly, each finding is assigned to a group, optionally selected from already generated group names. This section describes the different configurations explored for each step, including the version used in automatic timeline generation for annotation. We evaluate two LLMs, Llama 3.1 8B instruct (Grattafiori et al., 2024), referred to as Llama 3.1, and GPT-4o (Hurst et al., 2024).<sup>2</sup> We also evaluated GPT-OSS in preliminary

<sup>2</sup>The LLM prompts for each step are provided in Appendix B.

experiments but do not include it as we observed no performance gains over GPT-4o.

### 4.1. Step 1: Finding Extraction

We first extract lung-related findings from each report independently using an LLM. The prompt includes the full texts from a single report, detailed instructions about generating the lung findings, a short output example, and finally a request for the final answer. The instruction asks for details that may be relevant to risk evaluation, such as size and change. Because lung finding descriptions can be lengthy, we require the output format as a Python list in order to simplify post-processing.

When generating the timelines for annotation, we use Llama 3.1 in Step 1.

We automatically remove findings that contain only absent events according to an information extraction model, RadGraph (Delbrouck et al., 2022).<sup>3</sup> We use two versions of absent prediction filters. The first filters out sentences where all detected entities are absent. The second is more aggressive, dropping predictions when the core entities (e.g. lesions) are all absent even if their descriptors (e.g. lung) are not. For the 203 Llama-predicted findings, the first filter retains 187 findings, and the second reduces these to 165, closer to the 155 findings obtained in human labeling.

### 4.2. Step 2: Group Name Generation

Instead of directly clustering the findings, we use an intermediate step to generate interpretable group name in natural language first, which are used in the later group assignment step to map findings. By comparing findings with group names, we can avoid expensive pairwise finding comparison. Group names should be specific enough so that they can separate unrelated findings, and generic names should be avoided.

The prompts include longitudinal report records of a patient, detailed instructions asking for temporally related lung finding groups, one good (specific) and one bad (generic) group name example, and the final request asking for a list of group names separated by commas.

We explore two longitudinal input contexts to represent patient records; one includes the full reports, called **full context**, and the other includes only predicted findings from all reports in Step 1, called **finding-only context**. We use finding-only context to test whether filtering content into shorter lengths improves reasoning. In both input context versions,

<sup>3</sup>The RadGraph package is available at: <https://pypi.org/project/radgraph/>. Specifically, we use the model version `modern-radgraph-xl`.

the report contents are marked by dates and exam titles.

We use Llama 3.1 with full context in Step 2 to generate the timelines for annotation.

### 4.3. Step 3: Group Assignment

In this step, we assign each finding to a group name, so that findings with the same group name are grouped. We propose both LLM-based and more efficient embedding-based methods, each of which leverages the group names predicted in Step 2.

#### 4.3.1. LLM-based Approach

We explore prompting variations including single finding (**Single**) or multiple findings (**Multiple**) per LLM run, zero-shot (**ZS**) or few-shot (**FS**), and using full longitudinal contexts (**full context**) or **no context**.

For **Single** prompting, the LLM is prompted to assign each finding individually to the group best describing this finding among the provided group name list, or "Other" if no suitable group exists. The group name list is predicted in Step 2 for the same patient. If the prediction is 'Other', the finding's own text is used as its group name. The group names from Step 2 that are not assigned to any findings will be dropped from the final timelines. We instruct the LLM to place the selected group name in the last line of its response. During post-processing, the first group name appearing in the last line is parsed as the prediction, even if multiple group names are present.

When the generated group names are incomplete or duplicated, findings from the same group may not share consistent group names. To encourage consistent group assignment across all related findings, we experiment with **Multiple** prompting, which assigns groups to all findings of a patient in a single LLM run by inserting a tag next to each finding, following the approach of [Sundar et al. \(2024\)](#). Those tags serve as interpretable group names in natural language, and are parsed from angle brackets as the predicted group name for the corresponding finding. We use two variations. One does not require existing group names in the prompts (**no groups**), and the other requires them so that LLM can choose from existing group names or add new ones. This Multiple prompting is cheaper than Single prompting due to fewer runs.

We experiment with using all longitudinal radiology notes (**full context**) to provide complete cross-document information. This setup is similar to the human annotation setting where annotators have access to the full set of longitudinal reports. However, it remains unclear whether LLMs can effectively utilize such long contexts, and processing with full context is computationally expensive.

Therefore, we also evaluate **no context** prompts, where group assignment relies solely on the provided group names and the finding texts.

The few-shot (**FS**) prompts are motivated by formatting errors observed when zero-shot (**ZS**) prompts are used with Llama 3.1. We manually curate two short examples without using any examples from the evaluation dataset. Each example includes the existing group names if in use, and the output in the required format.

The final prompts integrate the above variations in the following order: longitudinal reports (omit for no context), detailed instructions (varies for Single or Multiple), two examples (omit for ZS), group names list (omit if no groups), a single or multiple findings to be grouped, and a final output request emphasizing on the output format.

We use Single ZS full-context with Llama 3.1 in Step 3 to generate timelines for annotation.

#### 4.3.2. Embedder-based Approach

We also explore embedding-based methods as more efficient alternatives for LLMs in the group assignment step. We first use a non-generative embedder model to convert each finding and each group name into embeddings, then assign each finding to the group name with the highest cosine similarity from the same patient. No new groups will be added beyond the provided group names.

Our embedder backbone is E5-Mistral which is built for general domain and supports instruction integration for task adaptation. We first describe the task using a **general instruction**, "Given a radiology finding, find the group that it belongs to", which is prepended to each finding. We further infuse task-specific preference using a **task-specific instruction**. We hypothesize that finding types (e.g., "opacities", "nodules") and anatomy locations (e.g., "right upper lobe") are critical information that distinguishes lung radiological groups. Therefore, we add a prefix, "Represent the radiological findings and be aware of the type and location of the findings", to both findings and group names.

## 5. Experiment and Results

We evaluate Llama 3.1 and GPT-4o for prompting methods in all steps. Context lengths of both are 128k tokens. We consistently choose greedy decoding for Llama 3.1 and set the temperature at 0.1, top-p at 0.9 for GPT-4o, without hyperparameter tuning. The group assignment step is evaluated on gold findings.

### 5.1. Finding Extraction

We compute ROUGE-L scores at both the report level and finding level, in both cases by first scor-

ing against each of the two annotators and then averaging the results. For the report-level score, we represent a report as the concatenated list of all findings before automatically removing absent findings, compute the ROUGE-L score for the list, and then average ROUGE-L scores across reports. When calculating the finding-level scores, we use the merged gold findings (findings from the same report assigned to the same group) and represent them using a concatenation of the respective finding strings. Then, we automatically align each predicted finding to the gold (merged) finding that yields the maximum ROUGE-L among all the gold findings from the same report, after filtering out absent findings. The finding-level ROUGE-L score is the average over the ROUGE-L scores for all matches.

Gold findings that have no match are considered **missed**. We reversely match each gold finding to the prediction that yields the maximum ROUGE-L and report the percentage of unmatched predictions as **unnecessary**. Under an ideal alignment, unnecessary predicted findings would include wrong, irrelevant, redundant, and unfiltered absent predictions.

Table 3 reports the different scores for finding extraction for two LLMs, Llama 3.1 and GPT-4o. Report-level ROUGE-L scores for Llama 3.1 and GPT-4o are close. For finding-level scoring, we provide two configurations, one with a weaker and one with a more aggressive absent finding filter. For both configurations, GPT-4o outperforms Llama 3.1 (67 versus 62). GPT-4o also has substantially lower unnecessary predictions. For examples, see Appendix C Table 16. Both LLMs have low rates of missed findings. As expected, the more aggressive filtering of absent findings leads to substantial reduction in the percent of unnecessary findings. In addition, the finding-level ROUGE-L score improves, with minimal increase in missed findings. The ideal alignment of reference to gold findings would yield roughly 3% missing and 39% unnecessary findings for Llama 3.1 predictions, which is close to the ROUGE-L alignment.

## 5.2. Group Name Generation

As timeline row headers, generated group names are useful for readability, but they also impact the later group assignment step as intermediate inputs.

Similarly to the finding-level ROUGE-L scores in Section 5.1, we compute the group-name-level ROUGE-L scores, missing and unnecessary prediction rates, as well as the percentage of gold names that have multiple predicted versions (duplicates) under automatic alignment. We compare the generated group names against two versions of gold group names separately, and report the average.

Model	RL (Rpt.)	RL (Fdg.)	Miss. %	Unnec. %
Llama-3.1	67	62	3	40
GPT-4o	68	67	5	25
Llama-3.1*	-	70	3	32
GPT-4o*	-	75	6	15

Table 3: Finding extraction evaluation against gold findings. ROUGE-L (RL) scores are at the report (Rpt.) and finding (Fdg.) level. Missing (Miss.) represents unmatched-gold percentage, and unnecessary (Unnec.) represents unmatched prediction percentage. The first two rows use the weaker absent finding filter; \* indicates to the more aggressive filter.

Table 4 reports the scores for the generated group names from two LLMs, using either full contexts or finding-only contexts, as well as the oracle scores evaluated between the two versions of gold group names. To calculate the oracle scores, we evaluate one gold group version against the other alternatively and take the averages. When using full contexts, GPT-4o has a higher ROUGE-L than Llama 3.1 (48 versus 38), but also a higher unmatched gold percentage (30% versus 25%). The Llama 3.1 model has nearly 60% unmatched predictions, suggesting considerably more unnecessary group names, likely due to irrelevant and duplicative predictions. When switching to finding-only contexts, the duplication issues in Llama 3.1 (from 32% to 42%) and unnecessary predictions in GPT-4o (from 37% to 46%) become more severe. In contrast, the missing percentage of GPT-4o decreases from 30% to 13%. These results suggest that using itemized findings as contexts produces a longer list of group names. Examples are given in Appendix C Table 17 and 18.

Method	RL	Miss. %	Dup. %	Unnec. %
Llama 3.1 (Full)	38	25	32	60
Llama 3.1 (Fdg.)	42	27	42	58
GPT-4o (Full)	48	30	28	37
GPT-4o (Fdg.)	52	13	31	46
Oracle	87	9	11	9

Table 4: Evaluation results for group name generation. The groups are generated using full or finding-only (Fdg.) contexts.

## 5.3. Group Assignment

We evaluate group assignment methods on gold findings, in order to directly compare with their as-

Group names	Single	Single	Multi.	Multi.	Single	Single	Multi.	Multi.	Embed.	Embed.
	ZS	FS	ZS	FS	ZS	FS	ZS	FS		
	Llama	Llama	Llama	Llama	GPT	GPT	GPT	GPT	Gen.	Spec.
Llama-3.1	66	74	77	78	77	73	78	77	68	70
GPT-4o	75	79	77	81	82	80	84	82	65	69
Oracle	73	78	86	84	80	84	85	83	78	83
No groups	n/a	n/a	72	68	n/a	n/a	70	72	n/a	n/a

Table 5: Group assignment performance of prompting and embedding-based methods with no context in CoNLL F1 scores. We use LLM models (Llama 3.1, GPT-4o) and an embedder model (e5-mistral-7b-instruct). "Multi." is short for Multiple prompting. "Embed. Gen." and "Embed. Spec." represent embedding-based methods using the general or task-specific instructions.

sociated gold group assignments. We consider findings that share the same group name and come from the same patient as a group, and report CoNLL F1 scores as in CDCR tasks. Given that we have two versions of gold group assignments, we calculate CoNLL F1 against each version and report the averages. We also report oracle results, assuming that gold group names are known and used as intermediate inputs, and further average the scores when evaluating against two gold group name versions separately.

### 5.3.1. Comparing Prompting and Embedder Methods without Contexts

Table 5 compares prompting methods and embedding-based methods when using no context in Step 3. We evaluate these methods when using generated group names, gold group names (oracle), or no group names. All generated group names are obtained from Step 2 using full contexts.

**The choice of intermediate group names is crucial, and GPT-4o approaches reach near-human performance on gold findings.** When shifting group names generated by Llama 3.1 to GPT-4o, all prompting methods improve, except that Multiple ZS prompt with Llama 3.1 is unchanged. The Single FS prompt with Llama 3.1 improves from 74 to 79 in CoNLL F1, and with GPT-4o it improves from 73 to 80. GPT-4o prompting methods with GPT-4o group names are all close to the inter-annotator performance at 82 CoNLL F1, and the best one, using Multiple ZS, reaches 84. Embedding-based methods do not show improvement with GPT-4o-generated names. Under oracle scenarios, prompting methods show mixed results compared to using automatically generated group names, indicating that automatic names are not necessarily worse than human curated ones. In contrast, embedding-based methods significantly improve with oracle names, and the best embedding-based method, which uses task-specific instructions, achieves near-human performance at 84 CoNLL F1, indicating their potential

as efficient substitutes to LLMs when group name generation is further improved. For the Multiple prompt version, we observe a considerable performance drop to around 70 after removing group names, suggesting the importance of including the group name generation step.

**Domain knowledge helps Embedding-based methods.** We compare the task-specific instruction and the general instruction for embedding-based methods. It shows that adding domain knowledge to instructions that emphasize the finding types and anatomy locations consistently boost results under all group name settings (from 78 to 83 with oracle names). The gap to LLM-based methods is narrower but still persists. More task adaptation efforts on embedding-based methods are needed.

**Few-shot prompting only helps Llama 3.1.** We observe that few-shot examples help Llama 3.1 in following a specified format. However, for GPT-4o, few-shot examples do not seem necessary.

### 5.3.2. Impact of Full Contexts

Context	Single	Single	Multi.	Multi.	Single	Multi.
	ZS	FS	ZS	FS	ZS	ZS
	Llama	Llama	Llama	Llama	GPT	GPT
Ann1						
No Ctx.	67	74	79	78	81	82
Full Ctx.	56	70	76	78	77	81
Ann2						
No Ctx.	65	75	75	77	82	85
Full Ctx.	60	76	82	78	86	90

Table 6: Impact of full contexts in group assignment prompting methods. Evaluated in CoNLL-F1 against gold group assignments from two annotators separately. Each configuration uses either Llama 3.1 for all steps or uses GPT-4o for all. All experiments use full contexts by default in group name generation. "Ctx." is short for "Context".

Table 6 shows the impact of adding full contexts for group assignment prompts, evaluated against two versions of gold group assignments separately. With the exception of Llama 3.1 Single ZS prompting, where full context hurts using both gold annotations, full context improves performance when measured against gold group assignments from annotator 2 but not annotator 1. These results suggest annotator differences in leveraging long-contexts, which is worth future investigation.

Group names	Single ZS	Single FS	Multi. ZS	Multi. FS
Llama (Full)	66	74	77	78
Llama (Fdg.)	66	70	71	68
GPT (Full)	82	80	84	82
GPT (Fdg.)	73	74	74	73

Table 7: Impact of input contexts for group name generation on the final group assignment performance in CoNLL-F1. All experiments use the same LLM in all steps and no contexts in group assignment. Group names are generated using full or finding-only ("Fdg.") contexts.

In Table 7, we further evaluate the impact of different group name generation prompts on group assignment. We observe that switching from full contexts to finding-only contexts for group name generation hurts the final group assignment performance, indicating that the additional context is useful.

## 6. Conclusion

We propose a timeline generation task for structured summarization on longitudinal radiology reports. This timeline format groups temporally related findings for straightforward comparison and facilitates fact-checking by denoting the associated report for each finding. We create a timeline dataset, RadTimeline, to evaluate timeline generation. We propose a three-step LLM approach that can achieve near-human performance in grouping gold findings.

This work can be extended to a broader clinical longitudinal records to capture temporal nuances in clinical narratives, such as for symptom progression, medicine history, and treatment trajectories.

## 7. Ethics

We receive approval from the Institutional Review Board (IRB) to access the patient data. All code and data are stored and executed on servers that complies with the Health Insurance Portability

and Accountability Act (HIPAA). We use a HIPAA-compliant version of APIs for GPT-4o experiments. All authors involved in running code and accessing the database have completed human subjects training. Our dataset will be de-identified to remove Protected Health Information (PHI) using an in-house tool (Lee et al., 2021). Release of the de-identified dataset is contingent on institutional approval.

## 8. Limitations

We report performance with only Llama 3.1 8B Instruct and GPT-4o. Due to data privacy concerns, our choice of proprietary LLMs is limited to those with HIPAA-compliant APIs.

When evaluating finding extraction and group name generation, we align predictions and gold references using a maximum ROUGE-L criterion. Using the ROUGE-based automatic alignment to score the gold findings give a rate of unnecessary predictions for Llama 3.1 in Table 3 that is consistent with aggregate statistics in the human annotation of those predictions. However, a more fine-grained analysis is needed to assess quality of the alignments.

Our method for evaluating grouping of findings requires that the findings have gold group assignments, which are only available for the gold findings. One solution would be to assign "silver" labels to the predicted findings by aligning them to gold findings, as in the finding assessment approach. Of course, this strategy should be validated through human evaluation of the fully automatic timeline.

The evaluation dataset contains only 10 patients, due to the costs associated with clinical excerpts and the cognitive effort required to understand longitudinal reports. This evaluation dataset gold annotations may share a similar language style to Llama 3.1, because the dataset is hand-corrected from generated timelines produced by Llama 3.1 based on expert feedback. This could lead to bias in ROUGE-L score.

We will de-identify the longitudinal reports when releasing the dataset. The performance may change after de-identification, and we will rerun our best methods to report performance on the de-identified dataset upon release.

## 9. Acknowledgment

The authors thank the annotators Emily Stiles, Nianiella Dorvall, Zixuan Yu for their contributions to this work. This work was supported in part by the National Institutes of Health (NIH) National Cancer Institute (NCI) (GrantNr. 1R01CA248422-01A1) and NIH Clinical and Translational Science Awards Program (CTSA) (GrantNr.UL1 TR002319, KL2

TR002317, and TL1 TR002318). The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

## 10. Bibliographical References

- Shany Barhom, Vered Shwartz, Alon Eirew, Michael Bugert, Nils Reimers, and Ido Dagan. 2019. [Revisiting joint modeling of cross-document entity and event coreference resolution](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4179–4189, Florence, Italy. Association for Computational Linguistics.
- Avi Caciularu, Arman Cohan, Iz Beltagy, Matthew Peters, Arie Cattan, and Ido Dagan. 2021. [CDLM: Cross-document language modeling](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2648–2662, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- A. Chien, H. Tang, B. Jagessar, K. W. Chang, N. Peng, K. Nael, and N. Salamon. 2024. [AI-assisted summarization of radiologic reports: Evaluating GPT3davinci, BARTcnn, LongT5booksum, LEDbooksum, LEDlegal, and LEDclinical](#). *AJNR American Journal of Neuroradiology*, 45(2):244–248.
- Surabhi Datta, Hio Cheng Lam, Atieh Pajouhi, Sunitha Mogalla, and Kirk Roberts. 2022. [A cross-document coreference dataset for longitudinal tracking across radiology reports](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3686–3695, Marseille, France. European Language Resources Association.
- Jean-Benoit Delbrouck, Pierre Chambon, Christian Bluethgen, Emily Tsai, Omar Almusa, and Curtis Langlotz. 2022. [Improving the factual correctness of radiology report generation with semantic rewards](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4348–4360, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Chen Huang and Guoxiu He. 2024. [Text clustering as classification with LLMs](#). *arXiv preprint, arXiv:2410.00927*.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Sameer Tajdin Khanna, Adam Dejl, Kibo Yoon, Steven QH Truong, Hanh Duong, Agustina Saenz, and Pranav Rajpurkar. 2023. [RadGraph2: Modeling disease progression in radiology reports via hierarchical information extraction](#). In *MLHC*, pages 381–402.
- K. Lee, N. J. Dobbins, B. McInnes, M. Yetisgen, and Ö. Uzuner. 2021. [Transferability of neural network clinical deidentification systems](#). *J Am Med Inform Assoc*.
- Tejas Sudharshan Mathai, Boah Kim, Oana M. Stroie, and Ronald M. Summers. 2025. [Privacy-preserving large language model for matching findings and tracking interval changes in longitudinal radiology reports](#). *Journal of Imaging Informatics in Medicine*. U.S. Government work; foreign copyright may apply.
- Qingkai Min, Qipeng Guo, Xiangkun Hu, Songfang Huang, Zheng Zhang, and Yue Zhang. 2024. [Synergetic event understanding: A collaborative approach to cross-document event coreference resolution with large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2985–3002, Bangkok, Thailand. Association for Computational Linguistics.
- Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang-Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [FACTScore: Fine-grained atomic evaluation of factual precision in long form text generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 12076–12100, Singapore. Association for Computational Linguistics.
- Niklas Muennighoff, Nouamane Tazi, Loic Magne, and Nils Reimers. 2023. [MTEB: Massive text embedding benchmark](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2014–2037, Dubrovnik, Croatia. Association for Computational Linguistics.
- Namu Park, Kevin Lybarger, Giridhar Kaushik Ramachandran, Spencer Lewis, Aashka Damani,

- Özlem Uzuner, Martin Gunn, and Meliha Yetisgen. 2024. [A novel corpus of annotated medical imaging reports and information extraction results using BERT-based language models](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 1280–1292, Torino, Italia. ELRA and ICCL.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. 2012. [CoNLL-2012 shared task: Modeling multilingual unrestricted coreference in OntoNotes](#). In *Joint Conference on EMNLP and CoNLL – Shared Task*, pages 1–40, Jeju Island, Korea. Association for Computational Linguistics.
- Kawshik Manikantan Sundar, Shubham Toshniwal, Makarand Tapaswi, and Vineet Gandhi. 2024. [Major entity identification: A generalizable alternative to coreference resolution](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11679–11695, Miami, Florida, USA. Association for Computational Linguistics.
- Dave Van Veen, Cara Van Uden, Maayane Attias, Anuj Pareek, Christian Bluethgen, Malgorzata Polacin, Wah Chiu, Jean-Benoit Delbrouck, Juan Zambrano Chaves, Curtis Langlotz, Akshay Chaudhari, and John Pauly. 2023. [RadAdapt: Radiology report summarization via lightweight domain adaptation of large language models](#). In *Proceedings of the 22nd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, pages 449–460, Toronto, Canada. Association for Computational Linguistics.
- Dave Van Veen, Cara Van Uden, Louis Blanke-meier, Jean-Benoit Delbrouck, Asad Aali, Christian Bluethgen, Anuj Pareek, Malgorzata Polacin, Eduardo Pontes Reis, Anna Seehofnerová, Nidhi Rohatgi, Poonam Hosamani, William Collins, Neera Ahuja, Curtis P. Langlotz, Jason Hom, Sergios Gatidis, John Pauly, and Akshay S. Chaudhari. 2024. [Adapted large language models can outperform medical experts in clinical text summarization](#). *Nature Medicine*, 30(7):1134–1142.
- R. Verma, E. Alsentzer, Z. Strasser, L. Chang, K. Roman, E. Gershanik, C. Hernandez, M. Linares, J. Rodriguez, D. Thakral, O. Unlu, J. You, L. Zhou, and D. Bates. 2025. [Verifiable summarization of electronic health records using large language models to support chart review](#). *medRxiv*. Preprint.
- Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. 2024. [Improving text embeddings with large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11897–11916, Bangkok, Thailand. Association for Computational Linguistics.
- Joonho Yang, Seunghyun Yoon, ByeongJeong Kim, and Hwanhee Lee. 2024. [FIZZ: Factual inconsistency detection by zoom-in summary and zoom-out document](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 30–45, Miami, Florida, USA. Association for Computational Linguistics.
- Fang Zeng, Zhiliang Lyu, Quanzheng Li, and Xiang Li. 2024. [Enhancing LLMs for impression generation in radiology reports through a multi-agent system](#). *arXiv preprint*, arXiv:2412.06828.
- Sihang Zeng, Yujuan Fu, Sitong Zhou, Zixuan Yu, Lucas Jing Liu, Jun Wen, Matthew Thompson, Ruth Etzioni, and Meliha Yetisgen. 2025. [Traj-CoA: Patient trajectory modeling via chain-of-agents for lung cancer risk prediction](#). In *The Second Workshop on GenAI for Health: Potential, Trust, and Policy Compliance*.
- Jin Zhao, Nianwen Xue, and Bonan Min. 2023. [Cross-document event coreference resolution: Instruct humans or instruct GPT?](#) In *Proceedings of the 27th Conference on Computational Natural Language Learning (CoNLL)*, pages 561–574, Singapore. Association for Computational Linguistics.

## Appendix

### A. Annotation Guidelines

The annotation pipeline consists of three sequential tasks. In the first task (Table 8), each cell is rated for factuality into levels of -1, 0, 1, 2. The second task (Table 9), collects final task group assignment on cells to form the timeline rows. Figure 2 shows one example presented to annotators. Finally, the third task (Table 10) asks an annotator to revise the group names on formed groups.

### B. Prompts

The prompts used in the three-step pipeline are provided in the subsequent tables. The step 1 prompt for finding extraction is shown in Table 11, and the step 2 group name generation prompt is presented in Table 12.

Step 3 group assignment prompts have three versions: single finding (Table 13), multi-finding (Table 15), and multi-finding without group names provided (Table 14).

### Task 1: Cell-level Factuality Rating

- For all cells in the column, rate the factuality level [0 bad, 1 partially correct, 2 good] for each cell based on the consistency with the column-associated report.
- Annotators should focus only on finding-related facts written in the specific report being evaluated. Do not refer to facts from other reports.
- Correctness of those facts should **NOT** affect the rating:
  - Suggestive conditions
    - \* e.g., “considering neoplasm given the persistent nodules”
    - \* e.g., “likely due to infection”
  - Facts that are not lung-related or not radiological findings
    - \* e.g., “Lung RAD rating is 3A”
    - \* e.g., “Coronary Artery Calcifications”
- **Skip ‘not-mentioned’ cells.** Blank cells are equivalent to ‘not mentioned’. Instead of annotating them, add them in Task 2 for missing findings if actually mentioned.
- Label the cell as not containing lung-related findings as **[-1: non-lung-findings]**.
- Add a note column next to the factuality cell.

Table 8: Annotation guidelines for factuality rating

### C. Qualitative Analysis

We present LLM generation examples for Step 1 finding extraction in Table 16 and Step 2 group name generation in Table 17 and Table 18. In both steps, GPT-4o generates fewer findings with greater conciseness, while Llama 3.1 tends to generate more findings and include content irrelevant to lung radiological findings. In Step 2, when switching from full longitudinal report context to the predicted finding list as a shorter context, GPT-4o produces more group names, possibly because the full context helps merge findings. However, this trend of reduced group names is not observed in Llama 3.1, potentially because the shorter context instead reduces irrelevant content in the Llama output.

### Task 2: Group Assignments

Assign a **group label** for each cell. The group labels to choose are the headers in the `finding_group` column. You may **add a new group label** if no existing group labels can distinguish this cell from other groups.

What should be included in a group

- The same finding in the notes, even if described differently, **or**
- Findings from different notes that can change to one another.

This may lead to trend descriptions such as:

- new
- increasing
- resolving

Duplicate labels If two group labels are duplicated, **choose one label and use it consistently** for all cells belonging to that group.

Goal:

Ensure that the same findings or time-related findings are placed in the same group.

Skip condition:

**Skip cells where factuality is 0 or -1.**

Table 9: Annotation guidelines for group assignments

### Task 3: Group Name Revision

Given a group of dated findings, add the group name for this group. The goal is to make the group name clearly distinguishable from other groups while remaining concise and containing critical information that a radiologist might use when referring to the finding.

Strategy:

The group name should:

1. **Include the facts shared by all members in the group.** When paraphrasing occurs, choose the description used more frequently.
2. **Omit details that can change in the future** (e.g., size, change trend, severity).
3. Use extra details if needed to distinguish similar groups.

Table 10: Annotation guidelines for group name revision given formed groups.

A	B	C	D	E	F
Rowid / group ID	Finding_group	[DATE1] _chest_ct	group ID	[DATE2] _chest_ct	group ID
1	Right upper lobe ground glass nodule	<b>10.2 mm ground glass nodule in the posterior segment of the right upper lobe with a small central nodular component</b>	1	<b>10 mm ground glass nodule in the right upper lobe with a central 4 mm cavity</b>	1
2	Right lower lobe ground glass nodule	2.0 cm ground glass opacity in the posterior segment of the right lower lobe with a small central nodular component	2	2 cm ground glass opacity in the posterior basal segment of the right lower lobe	2
3	Right upper lobe subsolid nodule	<b>2.6 mm noncalcified nodule in the anterior segment of the left upper lobe</b>	5	<b>4 mm subsolid nodule in the right upper lobe</b>	Skip if rated as wrong
4	Right lower lobe subsolid nodule				
5	Sub-3 mm nodule in the anterior upper lingula	Sub-3 mm nodule in the anterior upper lingula	Skip if rated as wrong	<b>Sub-3 mm nodule in the anterior upper lingula</b>	5
6	<b>Cardiomegaly</b>	<b>cardiomegaly</b>	Skip if rated as irrelevant		
7	<b>(new group name if need)</b>				

Figure 2: A group assignment example provided to annotators.

---

**Prompt for step 1 finding extraction**

---

**Q:** [Single Radiology Report Text, e.g. EXAMINATION: CT of the chest without contrast - low dose lung cancer screening..... ]:

Please list all lung-related findings (including pleura-related) from the exam report. Please use a python list to include the strings of those lung description. Each string should include all details related to higher or lower risk, using the same description in the note , especially finding's location, size, trend of change, and shape, density, or other characteristics.

Good example: ["stable 2 mm left upper lobe calcified nodule", "left lower lobe opacities"]

Now please answer:

lung\_finding\_list =

**A:** Here is the list of lung-related findings from the exam report:

“python lung\_finding\_list = [ "2.6 mm noncalcified nodule in the anterior segment of the left upper lobe", "alveolar groundglass opacity in the posterior segment of the right upper lobe with a small central nodular component measuring 10.2 mm", ...

---

Table 11: Step 1 finding extraction input. The prompt includes the single report input, an one-shot example, and instructions requiring a python list output format.

---

**Prompt for step 2 group name generation**

---

**Q:** [Longitudinal Textual Reports, e.g. date,note\_tag,QOClassificationName,ReportText YYYY-MM-DD,chest\_ct,CT Lung Low Dose Screening,"EXAMINATION: CT of the chest without contrast - low dose lung cancer screening ... ]:

Please group the lung-related findings recorded over time (including pleura-related findings) and list the group names.

Each group should contain findings that are essentially the same, even if described differently, as well as findings that evolve into one another over time.

Different findings or findings will not evolve into one another should be in separate groups

Group names should be specific that separate unrelated findings, avoid generic names.

Good group name: right upper lobe ground glass nodule

bad group name: ground glass nodules

Finally give me a list of group names separated by commas ','

**A:** .... List of group names: Right upper lobe ground glass nodule, Right lower lobe ground glass opacity with central nodular component, ...

---

Table 12: Step 2 group name generation prompt that requires list group names at the end of response. The prompt includes a good and a bad example for individual group names, but no examples for the list format.

---

**Prompt for step 3 group assignment for single finding**

---

**Q:** [Longitudinal Textual Reports, e.g. date,note\_tag,QOClassificationName,ReportText YYYY-MM-DD,chest\_ct,CT Lung Low Dose Screening,"EXAMINATION: CT of the chest without contrast - low dose lung cancer screening ... ]

Please find the group best describing this finding given the radiology exam history and group labels.

There are groups presented to track the lung-related findings recorded over time.

Each group should contain findings that are essentially the same, even if described differently, as well as findings that evolve into one another over time.

Different findings or findings will not evolve into one another should be in separate groups

**[Zero-Shot Content]**

The group labels are:

**[Group name list from step 2:** e.g. Right upper lobe ground glass nodule, Right lower lobe ground glass nodule, ...], other

The finding to be classified is: **[Single finding from step 1:** e.g. alveolar groundglass opacity in the posterior segment of the right upper lobe with a small central nodular component measuring 10.2 mm]

Give me the group in the last line starting with 'Answer:'.

**[Few-Shot Content]**

Here are some examples:

Example 1

Group labels: small RUL nodule, 10 mm right lower lobe nodule, other

Finding to be classified: Small 4 mm right upper lobe nodule stable

Answer: small RUL nodule

Example 2

Group labels: calcified granuloma, right upper lobe nodules, other

Finding to be classified: Benign appearing calcified granuloma in left lung

Answer: calcified\_granuloma

Your task:

Group labels: **[Group name list from step 2:** e.g. Right upper lobe ground glass nodule, Right lower lobe ground glass nodule, ...], other

Finding to be classified: **[Single finding from step 1:** e.g. alveolar groundglass opacity in the posterior segment of the right upper lobe with a small central nodular component measuring 10.2 mm]

Your explain: [EXPLAIN WHY]

Answer: [GROUP\_NAME CHOICE ]

Always end with "Answer: [GROUP\_NAME]" as the very last line

**A:** ... Answer: Right upper lobe ground glass nodule

---

Table 13: Step 3 Group Assignment for Single Finding. The prompts include underlined texts only when using full-contexts. We have both zero-shot and few-shot instructions, few-shot version adds two examples and further emphasize on the format requirements.

---

**Prompt for step 3 group assignment (multi-finding tag format, no group names given)**

---

**Q: Task:**

Your task is to group radiology findings evolving over time from multiple reports of the same patient.

Please replace every <#> with a group name tag in angle-bracket format, such as <10\_mm\_right\_lower\_lobe\_nodule>.

You should invent your own group tags so that each tag is representative for its associated group.

Use the same group name tag across all reports for all mentions related to the same underlying entity.

Your Task:

Radiology Findings:

**[Finding\_List\_with\_Markers,**

e.g. 2.6 mm noncalcified nodule in the anterior segment of the left upper lobe <#>

alveolar groundglass opacity in the posterior segment of the right upper lobe with a small central nodular component measuring 10.2 mm <#> ...]

Radiology Findings END

Please apply a group name tag consistently to all members of the group.

Please reply with exactly the same lines listed above that are between 'Radiology Findings' and 'Radiology Findings END', but with the group name tags inserted.

---

Table 14: Step 3 Group Assignment for Multi-Finding when no group names are given. The prompt includes an example of single tag, but no examples for the final output for all provided findings.

---

**Prompt for step 3 group assignment for multi finding**

**Q:** [Longitudinal\_Textual\_Reports, e.g. date,note\_tag,QOClassificationName,ReportText YYYY-MM-DD,chest\_ct,CT Lung Low Dose Screening,"EXAMINATION: CT of the chest without contrast - low dose lung cancer screening ... ]

Task: Your task is to group radiology findings evolving over time from multiple reports of the same patient. Those reports are as given above.

Please replace every <#> with a group name tag in angle-bracket format, such as <10\_mm\_right\_lower\_lobe\_nodule>. You can use either an existing group tag as provided, or invent your own.

Use the same group name tag across all reports for all mentions related to the same underlying entity.

**[Few-Shot Content]**

Here are examples:

**Example 1**

Existing group name tags:

<10\_mm\_right\_lower\_lobe\_nodule> : 10 mm right lower lobe nodule

<small\_RUL\_nodule> : small RUL nodule

Radiology Findings:

Small 4 mm right upper lobe nodule is new <#>

The 10 mm right lower lobe groundglass is stable <#>

Here are several left lobe calcified nodules <#>

Small right upper lobe nodule is stable <#>

RLL groundglass is smaller now <#>

Radiology Findings END

Answer:

Small 4 mm right upper lobe nodule is new <small\_RUL\_nodule>

The 10 mm right lower lobe groundglass is stable <10\_mm\_right\_lower\_lobe\_nodule>

Here are several left lobe calcified nodules <left\_lobe\_calcified\_nodules>

Small right upper lobe nodule is stable <small\_RUL\_nodule>

RLL groundglass is smaller now <10\_mm\_right\_lower\_lobe\_nodule>

**Example 2**

Existing group name tags:

<calcified\_granuloma> : calcified granuloma

<right\_upper\_lobe\_nodules> : right upper lobe nodules

Radiology Findings:

Benign appearing calcified granuloma in the left lung <#>

New right upper lobe nodules compared to last scan <#>

A 1.5 cm part-solid nodule is identified in the lingula <#>

Radiology Findings END

Answer:

Benign appearing calcified granuloma in the left lung <calcified\_granuloma>

New right upper lobe nodules compared to last scan <right\_upper\_lobe\_nodules>

A 1.5 cm part-solid nodule is identified in the lingula <part\_solid\_nodule\_in\_the\_lingula>

**[Few-Shot Content END]**

Your Task:

Existing group name tags:

[**Group name list from step 2:** e.g. right\_upper\_lobe\_ground\_glass\_nodule, right\_lower\_lobe\_ground\_glass\_nodule, ...]

Radiology Findings:

[**Finding\_List\_with\_Tags,**

e.g. 2.6 mm noncalcified nodule in the anterior segment of the left upper lobe <#>

alveolar groundglass opacity in the posterior segment of the right upper lobe with a small central nodular component measuring 10.2 mm <#> ...]

Radiology Findings END

If an existing group name tag fits, use it;

If multiple existing tags fit, choose one and apply it consistently to all members of the group.

If no existing group name fits, invent your own group tag.

Please reply with exactly the same lines listed above that are between 'Radiology Findings' and 'Radiology Findings END', but with the group name tags inserted.

Answer:

---

Table 15: Step 3 Group Assignment for Multi-Finding when group names are given. The prompts include underlined texts only when using full-contexts. We have both zero-shot and few-shot instructions, few-shot version adds two examples.

<b>Gold</b>	<b>GPT-4o</b>	<b>Llama3.1</b>
Moderate centrilobular emphysema stable lung nodules since [DATE], therefore benign. No new or enlarg- ing pulmonary nodule	Moderate centrilobular emphysema Stable lung nodules since [DATE], therefore benign  No new or enlarging pulmonary nod- ule	Moderate centrilobular emphysema stable lung nodules since [DATE], therefore benign. No new or enlarg- ing pulmonary nodule Normal supraclavicular region  Normal mediastinum Normal pleura Mildly dilated main pulmonary artery measuring 3.3 cm in caliber, also un- changed Stable mildly dilated pulmonary artery may represent pulmonary hy- pertension Stable mildly dilated ascending aorta may represent systemic hyperten- sion
Lungs: A few tiny (1–2 mm) nodules are scattered throughout the lungs [IMAGE NUMBES].	No large masses or dense airspace consolidations  A few tiny (1–2 mm) nodules are scat- tered throughout the lungs	Supraclavicular region: Normal.  Mediastinum: Normal.  Lungs: No large masses or dense airspace consolidations. Lungs: A few tiny (1–2 mm) nodules are scattered throughout the lungs. Pleura: Normal.
Low lung volume  Left base atelectasis	low lung volume with left base atelec- tasis no evidence of edema no pneumothorax	Cardiomegaly  Low lung volume Left base atelectasis No evidence of edema No pneumothorax
Right basilar consolidation likely pneumonia Patchy consolidation at the left lung base Left basilar consolidation likely pneu- monia Minimal density at right lung base	patchy consolidation at the left lung base minimal density at right lung base	Patchy consolidation at the left lung base Minimal density at right lung base  Right basilar consolidation likely pneumonia Left basilar consolidation likely pneu- monia
predominant fibrosis as before increased opacification of the left lower lobe	predominant fibrosis increased opacification of the left lower lobe best seen in the lateral view which may represent atelecta- sis, aspiration, or pneumonia	predominant fibrosis increased opacification of the left lower lobe  atelectasis, aspiration, or pneumonia (left lower lobe) no pleural effusion no pneumothorax

Table 16: Step 1 finding extraction examples from GPT-4o and Llama 3.1 8B Instruct. Llama 3.1 tends to generate more findings, and to include content irrelevant to lung findings. Human references are provided as the gold labels. The predicted findings are presented in the same order as they appear in the LLM responses.

Gold (group assignments from Ann1)	Gold (group assignments from Ann2)	GPT-4o	Llama3.1
<p>Hyperinflated lungs</p> <p>Flattening of domes of diaphragm</p> <p>Centrilobular emphysema in upper lobe predominant distribution</p> <p>Left anterior lingula scarring</p> <p>Spiculated nodule in left upper lobe (concerning for primary lung malignancy)</p> <p>Right upper lobe nodule</p> <p>Bilateral upper lobe lung nodules</p>	<p>Centrilobular emphysema in upper lobe predominant distribution</p> <p>Focal area of scarring</p> <p>Spiculated nodule in left upper lobe (concerning for primary lung malignancy)</p> <p>Right upper lobe nodule</p>	<p>Left upper lobe spiculated nodule</p> <p>Centrilobular emphysema</p> <p>Pleura and pleural spaces</p> <p>Pericardial effusion</p> <p>Ascending aorta abnormalities</p>	<p>Centrilobular emphysema in upper lobe predominant distribution</p> <p>Spiculated nodule in left upper lobe (concerning for primary lung malignancy)</p> <p>Unchanged spiculated left upper lobe nodule (6 x 9 mm)</p> <p>Unchanged right upper lobe nodule (2 mm)</p> <p>No new pulmonary nodules</p> <p>Pleural effusions or pneumothorax (no findings)</p> <p>Small bilateral fat-containing Bochdalek hernias</p> <p>Mild enlargement of the left lobe of the thyroid (supraclavicular region)</p> <p>Normal heart size</p> <p>Trace pericardial effusion (likely physiologic)</p> <p>Ascending aorta ectasia (unchanged)</p> <p>Mildly dilated main pulmonary artery (stable)</p> <p>Nonobstructive bilateral punctate renal calculi (upper abdomen)</p> <p>Moderate centrilobular emphysema</p> <p>Large lung volumes consistent with emphysema</p> <p>No acute bone or soft tissue abnormality</p>

Table 17: Step 2 group name generation examples from GPT-4o and Llama 3.1 8B Instruct when using full longitudinal reports as the context. Gold group names are provided in two versions, each reflecting the grouping assignments of Ann1 and Ann2 respectively, with names refined by a single third annotator across both versions. Similarly to step 1 results, Llama 3.1 tends to generate more group names and to include content irrelevant to lung findings. Two versions of human references are provided as the gold labels. The predicted group names are presented in the same order as they appear in the LLM responses.

<b>GPT-4o (Full)</b>	<b>GPT-4o (List)</b>	<b>Llama3.1 (Full)</b>	<b>Llama3.1 (List)</b>
Left upper lobe spiculated nodule	Hyperinflated lungs	Centrilobular emphysema in upper lobe predominant distribution	Hyperinflated lungs and diaphragm flattening
Centrilobular emphysema	Flattening of the domes of the diaphragm	Spiculated nodule in left upper lobe (concerning for primary lung malignancy)	Centrilobular emphysema
Pleura and pleural spaces	Centrilobular emphysema	Unchanged spiculated left upper lobe nodule (6 x 9 mm)	Spiculated nodule in left upper lobe (primary lung malignancy concern)
Pericardial effusion	Spiculated nodule in the left upper lobe	Unchanged right upper lobe nodule (2 mm)	Small focal area of scarring in lingula
Ascending aorta abnormalities	Right upper lobe pulmonary nodule	No new pulmonary nodules	Clear trachea and central airways
	Small focal area of scarring within the anterior aspect of the lingula	Pleural effusions or pneumothorax (no findings)	Pleural effusions or pneumothorax (absence of)
	No pleural effusions or pneumothorax	Small bilateral fat-containing Bochdalek hernias	Bilateral fat-containing Bochdalek hernias
	Small bilateral fat-containing Bochdalek hernias	Mild enlargement of the left lobe of the thyroid (supraclavicular region)	Large lung volumes consistent with emphysema
	Large lung volumes consistent with emphysema	Normal heart size	Stable lung nodules (benign)
	The lungs and pleural spaces are clear	Trace pericardial effusion (likely physiologic)	Normal supraclavicular region and mediastinum
		Ascending aorta ectasia (unchanged)	Normal pleura
		Mildly dilated main pulmonary artery (stable)	Mildly dilated main pulmonary artery (pulmonary hypertension concern)
		Nonobstructive bilateral punctate renal calculi (upper abdomen)	Mildly dilated ascending aorta (systemic hypertension concern)
		Moderate centrilobular emphysema	
		Large lung volumes consistent with emphysema	
		No acute bone or soft tissue abnormality	

Table 18: Comparison of step 2 group name generation examples using full longitudinal reports versus predicted finding lists as context. Examples are from GPT-4o and Llama 3.1 8B Instruct. The predicted group names are presented in the same order as they appear in the LLM responses.