

Reason2Decide: Rationale-Driven Multi-Task Learning

H M Quamran Hasan¹, Housam Khalifa Bashier¹, Jiayi Dai¹,
Mi-Young Kim², Randy Goebel¹

¹University of Alberta, Edmonton, Alberta, Canada

²University of Alberta, Camrose, Alberta, Canada

{hmquamra, khalifab, dai1, miyoung2, rgoebel}@ualberta.ca

Abstract

Despite the wide adoption of Large Language Models (LLMs), clinical decision support systems face a critical challenge: achieving high predictive accuracy while generating explanations aligned with those predictions. Current approaches suffer from exposure bias, leading to misaligned explanations. We propose Reason2Decide, a two-stage training framework that addresses key challenges in self-rationalization, including exposure bias and task separation. In Stage-1, our model is trained on rationale generation, while in Stage-2, we jointly train on label prediction and rationale generation, applying scheduled sampling to gradually transition from conditioning on gold labels to model predictions. We evaluate Reason2Decide on three medical datasets, including a proprietary triage dataset and public biomedical QA datasets. Across model sizes, Reason2Decide outperforms other fine-tuned baselines and some zero-shot LLMs in prediction (F1) and rationale fidelity (BERTScore, BLEU, LLM-as-a-Judge). In triage, Reason2Decide is rationale source-robust across LLM-generated, nurse-authored, and nurse-post-processed rationales. In our experiments, while using only LLM-generated rationales in Stage-1, Reason2Decide outperforms other fine-tuned variants. This indicates that LLM-generated rationales are suitable for pretraining models, reducing reliance on human annotations. Remarkably, Reason2Decide achieves these gains with models 40× smaller than contemporary foundation models, making clinical reasoning more accessible for resource-constrained deployments while still providing explainable decision support.

Keywords: Explainable AI, Explainability, Interpretability, Self-Explainable Models

1. Introduction

The integration of reasoning capabilities with prediction tasks has been a critical research problem in natural language processing (NLP). Existing state-of-the-art large language models (LLMs) struggle to balance high predictive accuracy while also generating human-interpretable explanations (Niu et al., 2025). Although LLMs have demonstrated strong performance on various question-answering benchmarks (Hendrycks et al., 2021; Srivastava et al., 2023), their ability to provide rationales that align with their predictions remains limited. This limitation is substantially amplified in healthcare where explanations are essential for trust and adoption.

Current approaches for rationale generation face two fundamental challenges. First, the exposure bias problem (Schmidt, 2019), which arises when models are trained to generate rationales conditioned only on ground-truth labels, but need to explain their own potentially incorrect predictions during inference. Second, the task separation problem, which occurs when prediction and rationale generation are treated independently rather than as mutually supporting tasks (Narang et al., 2020). While multi-task learning has shown promising performance, existing methods have failed to address the train-test discrepancy: models exclusively learn to explain gold-standard labels during training, leaving them unprepared for

rationalizing their own predictions during inference.

Recent work has acknowledged these challenges and explored various approaches to improve model interpretability. For example, “Chain-of-Thought” (CoT) prompting (Wei et al., 2022) demonstrates that step-by-step reasoning can improve complex reasoning; similarly, self-consistency techniques (Wang et al., 2023) show the value of aggregating multiple reasoning paths. However, these methods primarily focus on inference time strategies rather than addressing the fundamental training dynamics that help establish more robust reasoning capabilities. Knowledge distillation approaches (Hsieh et al., 2023) and rationale-augmented training (Lampinen et al., 2022) have also shown benefits, but do not train models to generate explanations conditioned on their own predictions during training.

Here we propose Reason2Decide (Figure 1), a framework which trains a single model to jointly predict and generate rationales. Our approach consists of a two-stage training regime, based on insights from curriculum learning (Bengio et al., 2009) and scheduled sampling (Bengio et al., 2015). The idea is to first learn to model explanation fundamentals and then jointly produce predictions and rationales, where we gradually shift from gold-label conditioning to self-conditioning. Through this process, the model learns to explain its own predictions. This addresses the exposure

bias problem while promoting alignment between predictions and rationales. The main contributions of this paper are:

- We propose Reason2Decide, a two-stage training framework for LLMs that first learns the fundamentals of explanation generation from rationales and then jointly optimizes predictions and rationales in a multi-task setup.
- We introduce a scheduled sampling mechanism that gradually transitions from gold labels to predicted label conditioning, mitigating exposure bias in self-rationalization.

2. Related Work

Our work builds upon research in rationale generation, multi-task learning, and existing methods to mitigate exposure bias.

2.1. Rationale Generation and Explainable AI

A primary goal of Explainable AI (XAI) is making model predictions transparent and interpretable. The development of early XAI methods focused on post-hoc explanation generation by analyzing learned model internal attributes or feature importance (Ribeiro et al., 2016; Lundberg and Lee, 2017). Recently, the field has shifted towards *self-explaining models* that can learn to generate rationales simultaneously with their predictions (Camburu et al., 2018). Our work aligns with the latter. For instance, Narang et al. (2020) fine-tuned T5 (Raffel et al., 2023) models to generate explanations by treating them as a separate task. Similarly, Lampinen et al. (2022) demonstrate that incorporating natural language explanations during training can improve reasoning and generalization. However, these approaches often treat the explanation as a secondary output, which creates a separation between the prediction and its explanation. In contrast, our Reason2Decide framework is designed to align the predictive task with the explanatory one, to help ensure that the explanation is an integral part of the reasoning process.

2.2. Multi-Task Learning and Knowledge Distillation

Multi-Task Learning aims to improve generalization by leveraging shared representations across related tasks (Ruder, 2017; Caruana, 1997). In NLP, models like T5 are considered multi-task learners because they combine diverse problems into a unified text-to-text format. We similarly use this foundation but specifically focus on the alignment between prediction and rationales. Closely

related work includes knowledge distillation approaches that use rationales. For example, Hsieh et al. (2023) distills the capabilities of a larger teacher model that generates step-by-step rationales into a smaller student model. While this approach can be effective in improving predictive accuracy, it does not explicitly enforce prediction–explanation coordination. The student model benefits from teacher rationales but is not trained to ensure that its own predictions and explanations are self-consistent. Our method addresses this weakness by explicitly training the coordination between prediction and rationale generation.

2.3. Mitigating Exposure Bias and Scheduled Sampling

Exposure bias in self-rationalization occurs when models are trained to generate explanations conditioned on gold labels but at inference time must justify their own predictions. Scheduled sampling (Bengio et al., 2015) addresses the issue by gradual replacement of teacher-forced tokens with model-generated ones. Here, we adapt this principle to the task level, thus treating the model’s predicted label as the conditioning context for rationale generation. Our mechanism systematically transitions from using gold labels to the model’s own predictions, directly mitigating this bias and preparing models for real-world deployment.

3. Methodology

Our method addresses the problem of combined prediction and rationale generation, where given an input x (clinical note or biomedical question), a model must predict a discrete label $y \in \mathcal{Y}$ and generate a free-text rationale r which justifies the prediction.

We employ a single encoder-decoder architecture f_θ (T5 variants) without modifications to its core components including the number of attention heads, hidden layer dimensions, relative positional encodings, and activation functions which are identical to the original specification. T5’s text-to-text formulation is ideal for our rationale-driven paradigm, as it naturally supports autoregressive decoding for both discrete label prediction (e.g., “Yes”, “No”, or “Go to ED Now.”) and free-form rationale generation from task-specific prompts. This allows seamless multi-tasking with explanations under a single objective, unlike encoder-only models like BERT (Devlin et al., 2019), which require separate classification heads and lack native support for conditional explanation generation.

Our training configuration is designed to handle both tasks: predict a discrete label and generate a free-text rationale. The model first predicts the la-

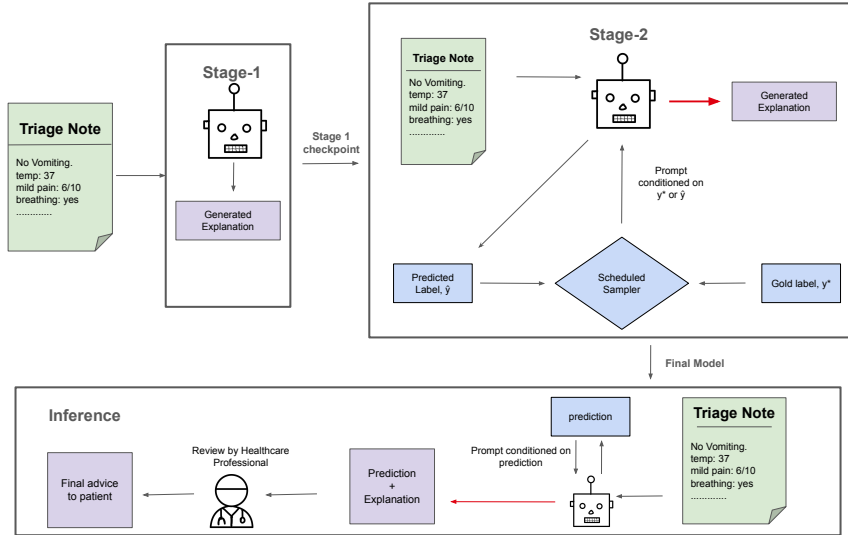


Figure 1: Overview of Reason2Decide. Stage-1 trains rationale generation. Stage-2 jointly predicts labels and generates label-conditioned explanations with task-level scheduled sampling. A single T5 model is used throughout; inference conditions explanations on the model’s own prediction.

bel as $\text{predict}: x \mapsto \hat{y}$, then employs the label to condition the rationale as: $\text{given label}: y^*$ or \hat{y} , $\text{explain}: x \mapsto \hat{r}$, where y^* is the gold label, \hat{y} is the model’s predicted label, and \hat{r} is the generated rationale.

3.1. Stage-1: Rationale Foundation Training

Previous work has shown that domain-specific pre-training significantly improves model performance in specialized tasks (Lewis et al., 2020; Lee et al., 2019). With this motivating background, our first stage teaches explanation fundamentals by training the model to generate rationales. Let θ denote the parameters of the T5 model. Given a gold rationale r^* , we optimize:

$$\mathcal{L}_{\text{stage1}} = -\log P_{\theta}(r^* | \text{explain}: x)$$

where $P_{\theta}(r^* | \cdot)$ denotes the autoregressive probability of the full rationale token sequence. The best model checkpoint is used to initialize Stage-2.

This two-stage structure follows a curriculum-inspired design, where the model first learns explanation fundamentals before being required to justify its own predictions.

3.2. Stage-2: Joint Optimization of Prediction and Explanation with Task-Level Scheduled Sampling

The second stage jointly optimizes prediction and explanation, starting with gold labels and gradually shifting to predicted labels for conditioning.

Prediction Task: The model treats label prediction as a text generation task, optimized via cross-entropy loss:

$$\mathcal{L}_{\text{pred}} = -\log P_{\theta}(y^* | \text{predict}: x)$$

where $P_{\theta}(y^* | \cdot)$ denotes the autoregressive probability of the full gold label sequence.

Explanation with Task-Level Scheduled Sampling:

To mitigate exposure bias in self-rationalization, we introduce a scheduled sampling mechanism. Unlike prior token-level scheduled sampling (Bengio et al., 2015), which replaces target tokens within sequences, we apply it at the task level by switching the conditioning context from gold labels to model-predicted labels. For instance, instead of gradually replacing reference tokens in a generated sequence, our method gradually replaces reference labels used for conditioning the rationale generation.

For each example we condition on either the gold label y^* or the model prediction \hat{y} (greedily decoded during training):

$$\tilde{y} = \begin{cases} y^* & \text{with probability } 1 - \pi_t \\ \hat{y} & \text{with probability } \pi_t \end{cases}$$

where π_t follows a linear schedule. With total training steps T , warm-up phase $w = 0.05T$, transition phase $m = 0.60T$, for step t :

$$\pi_t = \begin{cases} 0, & 0 \leq t < w, \\ \min\left(0.9, \frac{t-w}{m}\right), & w \leq t < w+m, \\ 0.9, & t \geq w+m. \end{cases}$$

The ceiling of 0.9 was a design choice to prevent the model from fully relying on self-generated labels during training, in order to help avoid error amplification from incorrect predictions. The hyperparameters (0.05 and 0.6) were determined through hyperparameter optimization using a few sets of predefined tuples, and are motivated by two key considerations:

Warm-up Phase: During $0 \leq t < w$, we linearly increase α_t from 0 to 0.7. This is to prioritize explanation loss $\mathcal{L}_{\text{expl}}$ over prediction loss $\mathcal{L}_{\text{pred}}$, to provide a smoother transition from single-task to multi-task optimization. After this phase α_t is held constant at 0.7 for the rest of the training. We heuristically selected 0.7 to balance both tasks, with a slight bias towards predictions, as Stage-1 already teaches rationale generation.

Transition Phase: Exposure bias arises in autoregressive rationale generation from the train-test mismatch: training conditions explanations on gold labels, but inference must use model predictions, risking error propagation. During $w \leq t < w + m$, we gradually increase the fraction of predicted labels used for conditioning from 0 to 0.9. This transition period is to systematically address exposure bias by training the model to learn coherent rationale generation conditioned on its own (potentially imperfect) predictions, bridging the gap between training and inference conditions.

The explanation loss is:

$$\mathcal{L}_{\text{expl}} = -\log P_{\theta}(r^* \mid \text{prompt})$$

where `prompt = given label: \tilde{y} , explain: x` . $P_{\theta}(r^* \mid \cdot)$ denotes the autoregressive probability of the full rationale token sequence.

The total loss combines both objectives with adaptive weighting:

$$\mathcal{L}_{\text{total}} = \alpha_t \mathcal{L}_{\text{pred}} + (1 - \alpha_t) \mathcal{L}_{\text{expl}}$$

3.3. Inference

During inference, we first predict the label via greedy decoding:

$$\hat{y} = \arg \max_y P_{\theta}(y \mid \text{predict}:x)$$

and then generate the rationale, via greedy decoding, conditioned on that prediction:

$$\hat{r} = \arg \max_r P_{\theta}(r \mid \text{prompt}),$$

where `prompt = given label: \hat{y} , explain: x` .

4. Experiments

Here we introduce the datasets and provide implementation details, followed by the experimental results.

4.1. Tasks and Datasets

We evaluate Reason2Decide on one proprietary clinical decision-making task (triage notes), and two public biomedical QA benchmarks (PubMedQA, BioASQ). As models we use T5-Small/Base/Large and zero-shot LLMs as non-fine-tuned references.

4.1.1. Clinical Triage Dataset

The Clinical Triage Dataset is derived from Alberta Health Link 811, a provincial telephone-based health advice service operated by Alberta Health Services (AHS), Canada. The dataset contains triage notes authored by nurses during telephone consultations with patients, along with the recommended care disposition and a rationale explaining the decision.

Task: Given a nurse triage note, predict a disposition (care pathway) and generate a rationale for that recommendation. The label space consists of 12 classes ranging from low (Home Care) to high urgency (e.g., "Go to Emergency Department (ED) now").¹

Rationale sources: We use three versions of rationales to assess source robustness:

- Nurse-authored rationales (original)
- Nurse post-processed rationales (lightly edited by an LLM)
- LLM-generated rationales (full generation using an LLM)

Table 1 shows sample rationales of each type. The post-processed version is there because the original version consists of atomic facts and is not always grammatically correct. We use² a Qwen-3 8B (Team, 2025) model to convert them into grammatically correct full sentences without adding new information. For LLM-generated notes, we provide the model the triage note and disposition and ask it to explain why the disposition was chosen. For experiments on the whole dataset, we used the LLM-generated rationales.

Data split: The train, validation and test sets consist of roughly 171k, 21k and 9.7k samples, respectively.

¹A complete list of the 12 disposition classes is provided in the Appendix, Subsection 10.2

²All prompts used are provided in the Appendix, Subsection 10.1.

Nurse-authored	Post-processed	LLM-generated
[1] Earache AND [2] MODERATE pain OR SEVERE pain inadequately treated per guideline advice - yes	The patient has an earache with moderate or severe pain inadequately treated according to guideline advice.	Persistent severe ear pain without fever or infection signs requires timely medical evaluation to prevent complications and ensure appropriate treatment.
[1] MILD-MODERATE pain AND [2] constant AND [3] present > 2 hours	The patient has mild to moderate pain that is constant and has been present for more than two hours.	Urgent evaluation needed due to persistent abdominal pain, bowel changes, and recent confusion, to rule out serious conditions and ensure appropriate treatment.

Table 1: Sample rationale variants.

4.1.2. PubMedQA

Task: Biomedical question answering with Yes/No/Maybe labels (Jin et al., 2019). We concatenate the question and context as the model input as triage note replacement.

Rationales: We use the dataset’s `long_answer` field as the gold-standard rationale.

Data split: We draw 100k stratified random samples (by labels) using a fixed seed from `pqa_artificial` (the dataset artificially generated), and create a 70k/20k/10k train/validation/test split. We then augment the test set with the 1k samples in `pqa_labeled`, as they are human annotated data, resulting in final splits of 70k (train), 20k(validation), and 11k (test: 10k `pqa_artificial` + 1k `pqa_labeled`).

4.1.3. BioASQ (Task B, Yes/No)

Task: Biomedical question answering dataset from the BioASQ 13 challenge (Nentidis et al., 2025). For our experiments, we focus on `yes/no` questions (binary classification).

Rationales: We extract and concatenate the `snippets` to serve as rationales for each question. Because the concatenated snippets can be

long, we first summarize them with a Qwen-3-8B model and use the summary as the gold rationale.

Data split: The official training dataset consists of 1,459 `yes/no` questions. We randomly select 250 questions for our validation set, and 1,209 questions for the train set. For the test set, we use the concatenation of the four official test batches, yielding 82 questions.

4.2. Implementation Details

We implement³ Reason2Decide using T5 Small (77M), Base (250M), Large (800M) architectures using A100 X 4 GPUs. Training uses AdamW optimizer (Loshchilov and Hutter, 2019) with learning rate 5×10^{-5} , max input length = 1024, and effective batch size = 64. Model-specific configuration is provided in Table 2.

For Stage-1, we monitor validation loss and early stop with patience of 3 validation evaluations. For Stage-2, we use delayed early stopping, with a patience of 5, that activates only after the scheduled sampling phase completes. This is to ensure the model fully benefits from our warm-up and scheduled sampling recipe.

During validation, we evaluate only the prediction task using F1-score, treating rationales as an auxiliary training signal. The best model checkpoint (highest F1-score) is loaded at the end of training. The models are trained using publicly available packages from <https://github.com/huggingface/transformers>.

Model	Per-GPU Batch	Grad Accum
T5-Small	16	1
T5-Base	4	4
T5-Large	2	8

Table 2: Training configuration by model size. Effective batch size is Per-GPU Batch \times #GPUs \times Grad Accum.

4.3. Baselines

We compare Reason2Decide against the following:

- **Standard fine-tuning (SFT):** Single-task label prediction without rationales (Howard and Ruder, 2018); model selection by validation Macro-F1.
- **Distilling Step-by-Step (DSS):** Multi-task training following (Hsieh et al., 2023). We run two selection criteria:

³Source code available at: <https://github.com/quamranhasan/Reason2Decide>

- **DSS-Loss:** Model selection by validation loss (as in Hsieh et al. (2023)).
- **DSS-F1:** Model selection by validation Macro-F1 (to align with our selection protocol).

Implementation follows the authors' release.

- **Zero-shot LLMs:** Zero-shot baselines without task-specific fine-tuning, using open-source models including Qwen-3-8B and Qwen-3-32B (Team, 2025), as well as Llama-3.1-Aloe-Beta-8B (Garcia-Gasulla et al., 2025).

Protocol: For all fine-tuned baselines, we use the same optimizer, learning rate, and effective batch size as in our method; per-device batch sizes are set following Table 2. Early stopping is used with patience of 5 validation evaluations. Model selection is by validation Macro-F1 unless otherwise stated (DSS-Loss). We report means over three seeds/runs. We use greedy decoding to ensure run-to-run determinism.

4.4. Evaluation

We provide a comprehensive evaluation across predictive performance and rationale quality:

Predictive Performance: We report Accuracy and Macro-F1 on the discrete label space (dispositions for triage; Yes/No/Maybe for PubMedQA; Yes/No for BioASQ). Scores are computed on held-out test sets. Due to the high class imbalance in the Triage Dataset (major and minor classes have 33, 633 and 239 samples respectively in the training set), we primarily rely on Macro-F1.

Rationale Fidelity: We assess explanation quality with:

- **BERTScore** (Zhang et al., 2020) (F1 variant): Semantic similarity between generated and gold rationales.
- **BLEU** (Papineni et al., 2002): N-gram overlap for surface-level quality.
- **LLM-as-a-Judge:** Following recent work demonstrating strong alignment between LLM and human evaluation (Zheng et al., 2023; Niu et al., 2025), we employ Qwen-3-8B for expert-style assessment on random 2k samples for the Triage Dataset and PubMedQA, and the whole test set for BioASQ (82 samples). We evaluate three defined metrics:

Correctness: 5-point scale for clinical alignment between the rationales and predictions.

The motivation is “If the predicted disposition is `Go to L&D now`, does the generated rationale justify the decision - not `homecare?`” To ensure consistent evaluation, we generated standardized disposition definitions using prompt-based refinement with a Qwen-3-32B model.

Coverage: 5-point scale that evaluates information retention. Scores measure how much relevant clinical information from the generated rationale is in the original triage note.

Overlap: 5-point scale for semantic content preservation between the gold and generated rationale.

For the triage dataset, we report all three LLM-as-a-Judge metrics. For PubMedQA and BioASQ, we only report `overlap`, as the others do not strictly align with QA datasets, due to missing triage notes and simple (`yes/no/maybe`) labels, respectively.

4.5. Results

4.5.1. Rationale Source Robustness

We evaluate Reason2Decide's robustness to rationale source variations by training Stage-1 exclusively on LLM-generated rationales, and in Stage-2, using the matching rationale variant (LLM-generated, nurse-authored, post-processed). As shown in Table 3, Reason2Decide maintains consistent performance across rationale variants while outperforming all baselines on both prediction and rationale metrics, with +0.37 to +6.09 F1 over DSS variants. Although this experiment was conducted on only 12% of the dataset, our method produces observable gains in data-scarce clinical situations.

Despite training on LLM-generated rationales only during Stage-1, Reason2Decide shows 51% lower F1 variation than DSS-F1 (1.08-point spread vs 2.21) and better rationale metrics. The F1 gains demonstrate that rationale pretraining provides transferable reasoning benefits that improve prediction, and not only explanation generation. This pattern suggests that LLM rationales may substitute for human-authored rationales during pretraining, thereby reducing reliance on costly human rationales. Our method demonstrates strong rationale source robustness, successfully adapting to different rationale styles during Stage-2 training, achieving best performance on post-processed rationales (51.14 F1) with consistent gains across rationale sources in BERTScore, BLEU, and LLM-as-a-Judge Overlap.

Rationale Source	Method	Macro F1	BERTSc.	BLEU	Overlap
	SFT	48.78 ± 0.44	-	-	-
LLM-generated	DSS-F1	49.69 ± 0.58	90.68 ± 0.12	15.00 ± 0.44	2.43
	DSS-Loss	47.27 ± 0.40	90.38 ± 0.05	13.72 ± 0.17	2.37
	Reason2Decide	50.06 ± 0.14	91.02 ± 0.02	16.46 ± 0.06	2.50
Nurse-authored	DSS-F1	48.01 ± 0.39	87.67 ± 0.06	24.88 ± 0.40	2.51
	DSS-Loss	46.76 ± 1.02	87.51 ± 0.07	23.87 ± 0.39	2.47
	Reason2Decide	50.76 ± 1.01	88.06 ± 0.23	26.62 ± 1.44	2.62
Post-processed	DSS-F1	47.48 ± 2.11	90.15 ± 0.13	22.47 ± 0.85	2.54
	DSS-Loss	45.05 ± 3.07	89.89 ± 0.03	20.68 ± 0.23	2.48
	Reason2Decide	51.14 ± 0.44	90.46 ± 0.07	24.03 ± 0.55	2.64

Table 3: Rationale Source Robustness on Clinical Triage (T5-Small on 12% of Clinical Triage Dataset)

Dataset	Model	Method	Macro F1	Accuracy	BERTSc.	BLEU
Clinical Triage	T5-Small	SFT	52.69 ± 0.86	-	-	-
		DSS-Loss	52.73 ± 0.99	-	91.14 ± 0.09	17.37 ± 0.45
		DSS-F1	52.09 ± 0.50	-	90.99 ± 0.06	16.72 ± 0.33
		Reason2Decide	55.88 ± 0.01	-	91.52 ± 0.06	19.33 ± 0.29
	T5-Base	SFT	57.10 ± 1.15	-	-	-
		DSS-Loss	53.26 ± 0.89	-	91.78 ± 0.04	20.86 ± 0.27
		DSS-F1	54.53 ± 2.63	-	91.56 ± 0.32	19.50 ± 1.70
		Reason2Decide	59.92 ± 0.42	-	92.09 ± 0.02	22.74 ± 0.06
	T5-Large	SFT	56.85 ± 7.21	-	-	-
		DSS-Loss	58.09 ± 0.77	-	92.08 ± 0.05	22.73 ± 0.17
		DSS-F1	59.43 ± 1.07	-	92.20 ± 0.03	23.56 ± 0.09
		Reason2Decide	60.58 ± 0.46	-	92.30 ± 0.03	24.13 ± 0.19
Aloe-8B	Zero-Shot	24.73 ± 0.00	-	-	-	
Qwen-3-8B		23.28 ± 0.00	-	-	-	
Qwen-3-32B		28.03 ± 0.00	-	-	-	
PubMedQA	T5-Small	SFT	52.18 ± 0.07	91.87 ± 0.34	-	-
		DSS-F1	53.63 ± 0.57	93.12 ± 0.16	89.14 ± 0.01	6.36 ± 0.06
		DSS-Loss	53.63 ± 0.07	93.14 ± 0.08	89.11 ± 0.05	6.27 ± 0.07
		Reason2Decide	55.02 ± 0.16	93.51 ± 0.14	89.30 ± 0.03	6.76 ± 0.15
	T5-Base	SFT	57.31 ± 0.16	94.58 ± 0.11	-	-
		DSS-F1	58.16 ± 0.11	95.00 ± 0.03	89.40 ± 0.04	6.93 ± 0.13
		DSS-Loss	57.69 ± 0.25	94.79 ± 0.14	89.40 ± 0.01	6.82 ± 0.10
		Reason2Decide	58.10 ± 0.22	95.04 ± 0.09	89.51 ± 0.03	7.17 ± 0.04
	T5-Large	SFT	59.60 ± 0.26	95.72 ± 0.17	-	-
		DSS-F1	59.92 ± 0.21	95.90 ± 0.06	89.62 ± 0.04	7.41 ± 0.12
		DSS-Loss	59.74 ± 0.26	95.70 ± 0.28	89.61 ± 0.03	7.08 ± 0.09
		Reason2Decide	60.28 ± 0.05	96.05 ± 0.01	89.60 ± 0.05	7.64 ± 0.15
Aloe-8B	Zero-Shot	45.33 ± 0.00	95.25 ± 0.00	-	-	
Qwen-3-8B		52.85 ± 0.00	87.89 ± 0.00	-	-	
Qwen-3-32B		58.33 ± 0.00	90.83 ± 0.00	-	-	
BioASQ	T5-Small	SFT	64.81 ± 2.09	72.76 ± 1.86	-	-
		DSS-F1	66.02 ± 1.19	71.14 ± 0.70	84.16 ± 0.43	2.68 ± 0.14
		DSS-Loss	40.47 ± 2.14	65.04 ± 1.41	82.80 ± 0.39	1.86 ± 0.15
		Reason2Decide	67.57 ± 2.26	73.58 ± 0.70	85.14 ± 0.18	3.27 ± 0.09
	T5-Base	SFT	53.70 ± 12.29	67.88 ± 2.54	-	-
		DSS-F1	66.28 ± 0.54	71.54 ± 1.86	85.78 ± 0.14	3.85 ± 0.02
		DSS-Loss	59.99 ± 13.58	71.14 ± 6.72	X	X
		Reason2Decide	68.02 ± 2.19	73.98 ± 1.41	85.64 ± 0.30	3.55 ± 0.10
	T5-Large	SFT	66.80 ± 2.64	73.98 ± 1.41	-	-
		DSS-F1	65.73 ± 2.15	70.73 ± 2.44	85.79 ± 0.25	3.43 ± 0.29
		DSS-Loss	44.55 ± 4.19	67.48 ± 1.41	X	X
		Reason2Decide	66.58 ± 4.57	73.17 ± 2.44	85.80 ± 0.11	3.39 ± 0.26
Aloe-8B	Zero-Shot	79.09 ± 0.00	81.71 ± 0.00	-	-	
Qwen-3-8B		50.32 ± 0.00	75.61 ± 0.00	-	-	
Qwen-3-32B		50.69 ± 0.00	76.83 ± 0.00	-	-	

Table 4: Performance comparison across datasets. Best F1 score per dataset is **bold+underlined**. Best fine-tuning strategy per model size is **bold**. Standard deviations are shown as ± values. Instances where rationale generation is non-applicable are shown as -. Instances where the model failed to generate a rationale are shown as X.

4.5.2. Consistent Performance Across Model Sizes and Tasks

From Table 4, across all three datasets, Reason2Decide shows improvement over other fine-tuned variants. For Clinical Triage, Reason2Decide outperforms every other baseline, across model sizes.

Similar results are observed with PubMedQA, with Reason2Decide outperforming other variants with the Small and Large models. Despite slightly lagging behind DSS-F1 with the Base model, Reason2Decide outperforms DSS-F1 on BLEU and BERTScore.

BioASQ proved to be a challenging dataset for all fine-tuned variants, possibly due to its limited data size. Although Reason2Decide outperforms the baselines on the Small and Base models, SFT takes the win with the Large model.

Reason2Decide outperformed both Qwen models on all datasets. However, Aloe-8B scores the highest on BioASQ. This can be attributed to its pretraining on biomedical literature, which closely matches the domain and format of BioASQ, but differs from the conversational and reasoning-oriented style of the triage and PubMedQA datasets.

While T5-Large (800M) is 40× smaller than the 32B Qwen model and is not a foundation model pretrained on medical corpora, it outperformed Qwen-3-32B on all three datasets and Aloe-8B on 2/3 datasets, narrowing the performance gap between LLMs and smaller language models.

4.5.3. Rationale Analysis

BLEU and BERTScore: As shown in Table 4, on Clinical Triage, for all three model sizes, Reason2Decide outperforms DSS variants in both metrics. On PubMedQA and BioASQ, we observe similar trends with Reason2Decide scoring competitively with DSS variants.

LLM-as-a-Judge Evaluation: To assess the quality of generated rationales, we evaluate three metrics as defined earlier. These were scored on a (1–5) Likert scale. From Table 5, Reason2Decide scores the highest in all metrics with the Small and Base model variants. With the Large model, the `overlap` and `coverage` scores are very close, with our proposed method scoring comparatively higher in `correctness`.

A sample prediction and rationale is provided in Table 6.

4.6. Ablation Study

We performed ablations to assess the contribution of each component in our Reason2Decide framework. Table 7 shows that removing Stage-1 consis-

Model Size	Method	Coverage	Correctness	Overlap
T5-Small	DSS-Loss	4.52	4.02	2.50
	DSS-F1	4.50	3.96	2.48
	Reason2Decide	4.62	4.26	2.58
T5-Base	DSS-Loss	4.74	4.09	2.57
	DSS-F1	4.71	4.05	2.55
	Reason2Decide	4.76	4.37	2.61
T5-Large	DSS-Loss	4.78	4.13	2.60
	DSS-F1	4.81	4.19	2.61
	Reason2Decide	4.80	4.43	2.63

Table 5: LLM-as-a-Judge evaluation metrics.

Method	Rationale and Prediction
Gold Standard	<i>Rationale:</i> Infant has repeated episodes of gasping for air, indicating possible serious breathing difficulty requiring immediate medical evaluation. <i>Prediction:</i> Go to ED Now
Reason2Decide	<i>Rationale:</i> Infant has repeated episodes of gasping for air, indicating possible serious breathing difficulty requiring immediate medical evaluation. <i>Prediction:</i> Go to ED Now
DSS-F1	<i>Rationale:</i> Infant's frequent gasping episodes with fever require immediate evaluation to rule out serious respiratory issues despite normal appearance and no worsening symptoms. <i>Prediction:</i> See Physician or PCP within 24 Hours
DSS-Loss	<i>Rationale:</i> Infant's frequent gasping episodes with normal breathing and no fever require prompt medical evaluation to rule out serious respiratory issues. <i>Prediction:</i> Go to ED Now

Table 6: Sample rationale and prediction. Reason2Decide produced gold-aligned rationale and prediction; DSS-F1 gave contradicting rationale and prediction; DSS-Loss made correct prediction but with inaccurate rationale.

tently hurts performance across datasets, confirming that explanation-focused pretraining improves downstream decision-making. Stage-2 is essential for predictions, because the Stage-1 model only learns to generate explanations.

The primary purpose of scheduled sampling is rationale alignment via exposure-bias mitigation. This is reflected with its equal or higher BERT and BLEU scores across both datasets. Additionally, it improves prediction on complex, multi-class datasets like Clinical Triage (59.92 vs. 57.28 F1), suggesting that training with realistic inference conditions could provide a greater benefit for challenging prediction tasks.

Including warm-up steps slightly improved performance on Clinical Triage across all metrics. Removing warm-up steps on BioASQ caused a substantial F1 drop despite small BERT/BLEU gains, indicating the importance of gradual adaptation from the single-task to the multi-task objective in Stage-2.

5. Conclusion and Future Work

We have introduced Reason2Decide, a two-stage training framework for LLMs designed to enhance decision quality and interpretability in clinical NLP tasks. Through this framework, the model learns to generate rationales that align with its predictions. Experiments on nurse triage and biomedical QA

Method	Accuracy	F1	BERT	BLEU
Clinical Triage				
w/o Stage-1	—	58.25	91.65	19.95
w/o Stage-2	—	0.00	91.78	20.72
w/o Scheduled sampling	—	57.28	92.09	22.59
w/o Warm-up steps	—	59.51	92.07	22.64
Reason2Decide	—	59.92	92.09	22.74
BioASQ				
w/o Stage-1	71.95	65.99	85.55	3.32
w/o Stage-2	0.00	0.00	84.17	2.95
w/o Scheduled sampling	72.76	68.06	84.15	3.07
w/o Warm-up steps	69.10	61.27	85.68	3.71
Reason2Decide	73.98	68.02	85.64	3.55

Table 7: Ablation results on Clinical Triage and BioASQ (Base model).

datasets show that Reason2Decide outperforms other fine-tuning variants in both prediction and explanation metrics. With our low cross-source F1 variation, we show that synthetic rationales can substitute for costly human rationales.

In future work, our aim is to extend this approach to broader clinical and biomedical subdomains, to assess its generalizability. We also want to investigate the effect of integrating rule-based or CoT-style rationales. In addition, our goal is to add human evaluation for better assessing the rationales.

6. Limitations

While Reason2Decide demonstrates strong performance over other fine-tuning variants, certain limitations remain. Firstly, the datasets used in this work fall under nurse triage and biomedical QA. To effectively assess Reason2Decide, it should be extended to other subdomains in clinical NLP. Secondly, the LLM-generated rationales were not rule-based. It would be valuable to see the impact of using CoT-style rationales during Stage-1 on predictions. Moreover, our two-stage training framework requires increased computational resources during training. However, inference is fully end-to-end from the final Stage-2 checkpoint alone: it both predicts the label and generates self-conditioned explanations. Our approach involved several hyperparameters that were chosen heuristically. A more granular hyperparameter optimization strategy may lead to Reason2Decide’s further improved performance. Finally, our rationale evaluation lacks human oversight. Although LLM-as-a-Judge methods are becoming increasingly adopted, they may be imperfect. For high-stakes domains like clinical NLP, human verification remains essential.

7. Ethics Statement

During this research, we ensured to follow ethical guidelines for clinical NLP. The proprietary dataset

used was de-identified. No personally identifiable information was accessible to the models or researchers during training/evaluation. Only open-source models were used, and were run on local machines. All clinical predictions and rationales generated by Reason2Decide should be treated as a decision support tool requiring human verification and not a replacement for human decisions.

8. Acknowledgements

We gratefully acknowledge Alberta Health Services (AHS) for providing access to the Alberta Health Link 811 dataset for research purposes. This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) Collaborative Research and Training Experience (CREATE) program “From Data to Decision” (FD2D). This research was also supported by the Alberta Machine Intelligence Institute (Amii), NSERC (including grants DGECR-2022-00369 and RGPIN-2022-0346), and Alberta Innovates.

9. Bibliographical References

- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled Sampling for Sequence Prediction with Recurrent Neural Networks. In *Proceedings of the 29th International Conference on Neural Information Processing Systems - Volume 1*, NIPS’15, page 1171–1179, Cambridge, MA, USA. MIT Press.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. [Curriculum learning](#). In *Proceedings of the 26th Annual International Conference on Machine Learning*, ICML ’09, page 41–48, New York, NY, USA. Association for Computing Machinery.
- Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. 2018. [e-SNLI: Natural Language Inference with Natural Language Explanations](#). In *Proceedings of the 32nd Conference on Neural Information Processing Systems*.
- Rich Caruana. 1997. [Multitask learning](#). *Mach. Learn.*, 28(1):41–75.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human*

- language technologies, volume 1*, pages 4171–4186.
- Dario Garcia-Gasulla, Jordi Bayarri-Planas, Ashwin Kumar Gururajan, Enrique Lopez-Cuena, Adrian Tormos, Daniel Hinjos, Pablo Bernabeu-Perez, Anna Arias-Duart, Pablo Agustin Martin-Torres, Marta Gonzalez-Mallo, et al. 2025. [The Aloe Family Recipe for Open and Specialized Healthcare LLMs](#). arXiv:2505.04388.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring Massive Multitask Language Understanding](#). arXiv:2009.03300.
- Jeremy Howard and Sebastian Ruder. 2018. [Universal Language Model Fine-tuning for Text Classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, Melbourne, Australia. Association for Computational Linguistics.
- Cheng-Yu Hsieh, Chun-Liang Li, Chih-kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alex Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. 2023. [Distilling Step-by-Step! Outperforming Larger Language Models with Less Training Data and Smaller Model Sizes](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8003–8017, Toronto, Canada. Association for Computational Linguistics.
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. [PubMedQA: A Dataset for Biomedical Research Question Answering](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2567–2577, Hong Kong, China. Association for Computational Linguistics.
- Andrew K. Lampinen, Nicholas A. Roy, Ishita Dasgupta, Stephanie C. Y. Chan, Allison C. Tam, James L. McClelland, Chen Yan, Adam Santoro, Neil C. Rabinowitz, Jane X. Wang, and Felix Hill. 2022. [Tell me why! Explanations support learning relational and causal structure](#). arXiv:2112.03753.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2019. [BioBERT: a pre-trained biomedical language representation model for biomedical text mining](#). *Bioinformatics*, 36(4):1234–1240.
- Patrick Lewis, Myle Ott, Jingfei Du, and Veselin Stoyanov. 2020. [Pretrained Language Models for Biomedical and Clinical Tasks: Understanding and Extending the State-of-the-Art](#). In *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, pages 146–157, Online. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled Weight Decay Regularization](#). arXiv:1711.05101.
- Scott Lundberg and Su-In Lee. 2017. [A Unified Approach to Interpreting Model Predictions](#). In *Advances in neural information processing systems*.
- Sharan Narang, Colin Raffel, Katherine Lee, Adam Roberts, Noah Fiedel, and Karishma Malkan. 2020. [WT5?! Training Text-to-Text Models to Explain their Predictions](#). arXiv:2004.14546.
- Anastasios Nentidis, Georgios Katsimpras, Anastasia Krithara, Martin Krallinger, Miguel Rodríguez-Ortega, Eduard Rodríguez-López, Natalia Loukachevitch, Andrey Sakhovskiy, Elena Tutubalina, Dimitris Dimitriadis, Grigorios Tsoumakas, George Giannakoulas, Alexandra Bekiaridou, Athanasios Samaras, Giorgio Maria Di Nunzio, Nicola Ferro, Stefano Marchesin, Marco Martinelli, Gianmaria Silvello, and Georgios Paliouras. 2025. [Overview of BioASQ 2025: The Thirteenth BioASQ Challenge on Large-scale Biomedical Semantic Indexing and Question Answering](#). In *Proceedings of the International Conference of the Cross-Language Evaluation Forum for European Languages*, pages 173–198.
- Shuai Niu, Jing Ma, Hongzhan Lin, Liang Bai, Zhihua Wang, Richard Yi Da Xu, Yunya Song, and Xian Yang. 2025. [Knowledge-Augmented Multimodal Clinical Rationale Generation for Disease Diagnosis with Small Language Models](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11011–11024, Vienna, Austria. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [BLEU: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2023. [Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.

Marco Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. “why should I trust you?”: Explaining the predictions of any classifier. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, pages 97–101, San Diego, California. Association for Computational Linguistics.

Sebastian Ruder. 2017. [An Overview of Multi-Task Learning in Deep Neural Networks](#). arXiv:1706.05098.

Florian Schmidt. 2019. [Generalization in Generation: A closer look at Exposure Bias](#). In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pages 157–167, Hong Kong. Association for Computational Linguistics.

Aarohi Srivastava et al. 2023. [Beyond the Imitation Game: Quantifying and extrapolating the capabilities of language models](#). *Transactions on Machine Learning Research*.

Qwen Team. 2025. [Qwen3 Technical Report](#). arXiv:2505.09388.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-Consistency Improves Chain of Thought Reasoning in Language Models](#). In *Proceedings of The Eleventh International Conference on Learning Representations*.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. [Chain-of-Thought Prompting Elicits Reasoning in Large Language Models](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [BERTScore: Evaluating Text Generation with BERT](#). In *Proceedings of the International Conference on Learning Representations*.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena](#). In *Advances in neural information processing systems*, pages 46595–46623.

10. Appendix

10.1. Prompts Used

Prompt for Post-Processing Nurse-Authored Rationales:

You are a helpful assistant who expands brief medical notes into full, grammatically correct sentences using fewer than 20 words. Do not add new information.
Convert this to a sentence without adding new information: [RATIONALE]

This created the nurse post-processed rationale variant by replacing [RATIONALE] with each original nurse note.

Prompt for Generating Standardized Disposition Definitions:

You are building a clinical definition for the disposition: [DISPOSITION_NAME].

Current working definition:
[CURRENT_DEFINITION]

You are given new examples of triage notes and rationales for this disposition. Use them to refine, expand, or correct the working definition. Keep the definition concise but clinically accurate.

Triage Notes and Rationales: [EXAMPLES]

Update the definition:

- Incorporate any new key symptoms, criteria, or thresholds.
- Remove incorrect parts.
- Keep it as clear and specific as possible.

Output only the revised definition text, nothing else.

Where [DISPOSITION_NAME], [CURRENT_DEFINITION], and [EXAMPLES] were replaced with actual data. This process generated the standardized definitions used for consistent LLM-as-a-Judge scoring.

Prompt for Zero-Shot Disposition Classification:

Issue: [ISSUE_ASSESSMENT]
Dispositions: [CLASSES_TEXT]

Classify the healthcare issue into one of the dispositions above. Return your answer in the following ****strict**** format:
Class: [chosen digit]

Do not ask for more information, and do not provide any general statements. Only respond with the digit.

Where [ISSUE_ASSESSMENT] contained the triage note text and [CLASSES_TEXT] listed the 12 disposition options with their numerical identifiers. This prompt was used for zero-shot LLM evaluation.

Prompt for Summarizing BioASQ Snippets:

Please summarize the following medical rationale in 100 words or less. Focus on the key points and main conclusions. Keep it concise and informative.

Rationale to summarize: [RATIONALE]
Summary:

Where [RATIONALE] was replaced with the concatenated text snippets from BioASQ. This created concise gold-standard rationales for BioASQ experiments.

Prompt for LLM-as-a-Judge Correctness Scoring:

You are an expert clinical trainer for telephone triage nursing.
[DEFINITIONS]

Task: Given the following rationale and disposition, score the alignment on a scale of 1 to 5, where:

- 5 - Excellent Alignment
- 4 - Good Alignment
- 3 - Moderate Alignment
- 2 - Poor Alignment
- 1 - Very Poor Alignment

Rationale: [RATIONALE]
Disposition: [DISPOSITION]
Output exactly one number (1, 2, 3, 4, or 5) with no other text.

Where [DEFINITIONS] was replaced with the standardized disposition definitions following Subsection 10.2, [RATIONALE] with the generated rationale, and [DISPOSITION] with the predicted disposition.

Prompt for LLM-as-a-Judge Coverage Scoring:

You are a medical text comparison assistant.
Task: Compare the triage note with the text provided. Determine how much of the information in the text is also present in the triage note.

Output only one number based on this scale:
5 = The note contains everything mentioned in the text.
4 = The note contains most things mentioned in the text.
3 = The note contains some but not most things mentioned in the text.
2 = The note contains almost nothing mentioned in the text.
1 = The note contains nothing mentioned in the text.

Triage Note: [NOTE]
Text: [TEXT]
Output only the number:

Where [NOTE] was replaced with the original triage note and [TEXT] with the generated rationale.

Prompt for LLM-as-a-Judge Overlap Scoring:

You are a clinical text comparison assistant.
Task: Compare the nurse's original rationale with the model's predicted rationale and rate how much of the nurse's rationale is present in the predicted rationale. Judge semantic

content, not wording. Paraphrases/synonyms count as overlap. Ignore pleasantries and generic instructions. Contradictions reduce overlap.

Output exactly one number per line for each pair, with no other text:

- 1 = No overlap
- 2 = Almost no overlap
- 3 = Some but not most overlap
- 4 = Most overlap
- 5 = Complete overlap

Nurse Rationale: [NURSE_RATIONALE]
Predicted Rationale:
[PREDICTED_RATIONALE]
Output only the number:

Where [NURSE_RATIONALE] was the gold-standard nurse rationale and [PREDICTED_RATIONALE] was the model-generated rationale.

10.2. Clinical Disposition Definitions

Home Care: Patients suitable for Home Care disposition have mild to moderate, stable, and non-progressive symptoms without signs of immediate or severe complications requiring emergency intervention. They exhibit no respiratory distress, hemodynamic instability, severe pain unresponsive to treatment, significant bleeding, high fever with systemic symptoms, spreading infection, altered mental status, neurological deficits, or other urgent clinical concerns. Their condition allows for safe management, symptom relief, and observation at home with appropriate follow-up. This includes stable minor injuries, controlled localized infections, mild chronic condition issues, and non-urgent questions or concerns. Patients and caregivers should be advised to seek urgent care if symptoms worsen or new concerning signs develop.

See Physician or PCP within 3 days: Patients with mild to moderate, stable but persistent or worsening symptoms that do not pose an immediate threat to life or function, and who require timely medical evaluation to prevent complications. This includes conditions without severe pain, respiratory distress, hemodynamic instability, significant neurological deficits, acute infection, or other urgent signs. The disposition excludes any patients exhibiting signs of severe or rapidly progressing illness, acute psychiatric emergencies, or other urgent conditions necessitating immediate or emergency care.

Call Pharmacist within 24 Hours: This disposition is used when a caller has non-emergency medication-related questions or concerns that require timely pharmacist expertise within 24 hours to ensure safe, effective, and appropriate medication use. It applies when immediate urgent care is not needed but professional assessment is necessary to clarify dosing, manage side effects, verify medication information, address potential interactions, support adherence, assist with medication access, or provide guidance on special populations and circumstances. This ensures optimized therapy, prevention of harm, and informed patient decisions without delay.

See Physician or PCP within 4 Hours (or PCP triage): Patients with new or worsening moderate to severe symptoms, signs of infection, or conditions at risk of rapid deterioration that require timely clinical evaluation within hours to prevent complications. This includes significant pain unrelieved by initial treatment, progressive neurological symptoms, signs of systemic infection, post-procedural complications, unstable vital signs, moderate to severe respiratory,

abdominal, or urinary symptoms, pregnancy-related concerns, dehydration, metabolic instability, mental health deterioration, and other clinical presentations indicating potential for rapid decline. The disposition ensures prompt physician assessment to guide urgent management and prevent adverse outcomes.

Call EMS 911 Now: Initiate immediate emergency medical services activation for any patient presenting with signs or symptoms of a potentially life-threatening condition. This includes acute neurological deficits, severe chest pain or cardiac symptoms, significant respiratory distress, signs of anaphylaxis, major trauma or uncontrolled bleeding, severe abdominal or back pain with systemic symptoms, acute deterioration in patients with serious underlying conditions, active suicidal intent with risk of harm, critical illness in infants or young children, shock or imminent collapse, severe infection or sepsis, altered mental status or unresponsiveness, and any situation posing an immediate threat to life, safety, or vital functions. Prompt EMS activation is essential whenever there is concern for compromised airway, breathing, circulation, neurological status, or urgent mental health crisis requiring rapid intervention.

Go to ED Now: Immediate emergency department evaluation is required for patients presenting with sudden, severe, or rapidly worsening symptoms that pose an immediate risk to life, limb, or function. This includes active uncontrolled bleeding, significant head or neck injuries, severe respiratory distress or airway compromise, acute neurological deficits, severe or persistent pain suggestive of surgical or obstetric emergencies, signs of shock or altered consciousness, severe allergic reactions, prolonged seizures, suspected serious infections, severe metabolic disturbances, and any other critical conditions requiring urgent assessment and intervention.

See Physician or PCP within 2 Weeks: Patients with new, persistent, worsening, or recurrent symptoms that are stable and do not require emergency care but need timely medical evaluation to diagnose, monitor, or adjust treatment. This includes conditions without signs of acute distress, hemodynamic instability, severe pain, neurological deficits, respiratory compromise, systemic infection, or other urgent symptoms. The disposition applies to a broad range of non-emergent but concerning clinical presentations where prompt follow-up is necessary to prevent progression or complications.

Go to L&D Now: Immediate evaluation in Labor

and Delivery is required for pregnant individuals presenting with signs of active labor at or near term, suspected or confirmed rupture of membranes, significant vaginal bleeding, decreased or absent fetal movement, new or worsening moderate to severe abdominal or pelvic pain, signs of preterm labor, pregnancy complications or risk factors combined with concerning symptoms, abdominal or pelvic trauma, maternal conditions suggestive of serious complications (e.g., preeclampsia, infection, hemodynamic instability), or any other acute symptoms indicating potential maternal or fetal compromise. Prompt assessment is essential to ensure maternal and fetal safety.

Call Poison Center Now: Immediate expert toxicology consultation is required for any suspected or confirmed exposure to potentially harmful substances or situations with risk of significant toxicity, overdose, or complications. This includes exposures involving high-risk medications, chemicals, toxins, unknown or unlabeled agents, vulnerable populations (such as children or pregnant women), or any new or worsening symptoms suggestive of systemic toxicity. Prompt specialist guidance is essential to ensure safe management and appropriate treatment.

See More Appropriate Guideline: Use this disposition when the caller's concerns do not require emergency or urgent care but need assessment, advice, or management under a more specific, condition-focused guideline. It applies to stable, non-urgent symptoms or questions without red flags, including chronic conditions, mild new symptoms, or informational and care coordination needs. Avoid this disposition for any signs of acute deterioration, emergencies, or conditions requiring immediate intervention. Direct callers to the guideline that best matches their specific symptom or concern to ensure appropriate care.

See Physician or PCP within 24 Hours: Patients with new, worsening, or persistent moderate symptoms that impact daily activities but do not require emergency care. This includes localized signs of infection or inflammation without systemic involvement, moderate injuries without severe complications, mild to moderate neurological, respiratory, gastrointestinal, or mental health symptoms without acute distress, and other conditions needing timely medical evaluation to prevent deterioration, ensure appropriate management, and monitor progression.

Call Dentist when Office is Open: This disposition applies to patients with non-emergent dental issues that do not exhibit signs of serious

infection, airway compromise, uncontrolled bleeding, or systemic illness. It includes mild to moderate pain or discomfort manageable with analgesics, stable post-procedure symptoms, minor dental trauma without active bleeding or severe pain, and dental appliance-related discomfort without urgent complications. Patients should seek dental care during regular office hours and be advised to obtain immediate emergency care if symptoms worsen or signs of a dental or medical emergency develop.