

Overview of the First Shared Task on Clinical Text Generation: RRG24 and “Discharge Me!”

Justin Xu^{*1}

Zhihong Chen^{*1}

Andrew Johnston¹

Louis Blankemeier¹

Maya Varma¹

Jason Hom¹

William J. Collins¹

Ankit Modi²

Robert Lloyd²

Benjamin Hopkins³

Curtis P. Langlotz¹

Jean-Benoit Delbrouck¹

¹Stanford University ²University of Arizona

³University of Southern California

{xujustin,zhihongc,drewj32,langlotz,jbdel}@stanford.edu

Abstract

Recent developments in natural language generation have tremendous implications for healthcare. For instance, state-of-the-art systems could automate the generation of sections in clinical reports to alleviate physician workload and streamline hospital documentation. To explore these applications, we present a shared task consisting of two subtasks: (1) Radiology Report Generation (RRG24) and (2) Discharge Summary Generation (“Discharge Me!”). RRG24 involves generating the ‘Findings’ and ‘Impression’ sections of radiology reports given chest X-rays. “Discharge Me!” involves generating the ‘Brief Hospital Course’ and ‘Discharge Instructions’ sections of discharge summaries for patients admitted through the emergency department. “Discharge Me!” submissions were subsequently reviewed by a team of clinicians. Both tasks emphasize the goal of reducing clinician burnout and repetitive workloads by generating documentation. We received 201 submissions from across 8 teams for RRG24, and 211 submissions from across 16 teams for “Discharge Me!”.

1 Introduction

An important application of natural language generation (NLG) in medical artificial intelligence (AI) is radiology report generation (RRG). Specifically, an RRG system can be designed to accept radiology images (*e.g.*,

chest X-rays) of a patient and generate a textual report describing the clinical observations in the images. This is a clinically important task, and offers the potential to reduce the repetitive work of radiologists and generally improve clinical communication (Pang et al., 2023). Existing studies have been conducted using a single dataset, which limits the scale and diversity of the data and results. Therefore, we introduce our first subtask, RRG24, where we curate Interpret-CXR, a large-scale collection of RRG datasets from a variety of different sources (*i.e.*, MIMIC-CXR (Johnson et al., 2019), CheXpert (Irvin et al., 2019), PadChest (Bustos et al., 2020), BIMCV-COVID19 (Vayá et al., 2020), and OpenI (Demner-Fushman et al., 2016)). In RRG24, participants generate the Findings and Impression sections from chest X-rays. We then evaluate the generations on common leaderboards with standard and recently proposed metrics. Ultimately, this task aims to benchmark recent progress using common data splits and evaluation implementations.

NLG can also impact discharge documentation by playing a role in generating discharge summaries. Hence, we introduce our second subtask, “Discharge Me!”, with the primary objective of encouraging NLG systems that alleviate clinician burden when writing detailed discharge summaries. Clinicians play a crucial role in documenting patient progress after a hospital stay, but the creation of concise yet comprehensive Brief Hospital Course (BHC) sections and Discharge Instructions often demands a significant investment of time (Do et al., 2020; Alissa et al., 2021). These two sections in particular cannot be readily copied from prior notes, and thus must be written from scratch by clinicians who synthesize information from across the patient record (Weetman et al., 2021). This process

^{*}Equal contribution

contributes to clinician burnout and poses operational inefficiencies within hospital workflows (Haycock et al., 2014). We hypothesize that computer-generated clinical documentation has the potential to more accurately and completely capture a patient’s hospital course while reducing the administrative burden on clinicians, which, in turn, mitigates burnout, streamlines hospital operations, and ultimately improves the quality of care. Thus, in “Discharge Me!”, participants submit generations of both target sections (BHC & Discharge Instructions). We evaluate submissions on a common leaderboard and conduct a subsequent manual clinician review to measure clinical alignment of the outputs.

2 Related Work

2.1 Radiology Report Generation

Recent advances in computer vision (CV) and NLG have shown great potential for the automatic generation of radiology reports. This progress can be summarized from three perspectives:

- (1) Data: Most relevant studies focus on chest X-rays, mainly owing to the current number of publicly available image-report datasets for this modality (e.g., MIMIC-CXR, PadChest, and OpenI, etc.). Recently, there have also been studies expanding the scope of radiology report generation to other modalities (e.g., computed tomography (CT) (Loveymi et al., 2021; Hamamci et al., 2024) and ultrasound (Zeng et al., 2020; Yang et al., 2021; Huh et al., 2023)).
- (2) Methodology: The methods for radiology report generation have evolved from task-specific modeling to pre-training-based approaches. For the former, researchers have incorporated the task priors into the designs of the model architectures (Shin et al., 2016; Zhang et al., 2017; Jing et al., 2018; Chen et al., 2020; Zhang et al., 2020; Liu et al., 2021; Delbrouck et al., 2022a; Hou et al., 2023), whereas for the latter, researchers have performed domain-specific representation learning using vision encoders or have adopted large pre-trained language decoders (Thawkar et al., 2023; Hyland et al., 2023; Tu et al., 2024).
- (3) Evaluation: One of the largest factors hampering radiology report generation progress is the selection of evaluation metrics. Due to its domain-specific characteristics, simple n -gram matching metrics (e.g., BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), METEOR (Banerjee and Lavie, 2005), and CIDEr (Vedantam et al., 2015)) are sub-optimal choices for this task. However, researchers have proposed various model-based metrics for evaluating the quality of generated reports, such as BERTScore (Zhang et al., 2019), F1-CheXbert (Smit et al., 2020), F1-RadGraph (Delbrouck et al., 2022a), and GREEN (Ostmeier et al., 2024).

2.2 Discharge Summary Generation

Previous research has also examined AI technologies for the generation of discharge summaries to alleviate clerical burden for clinicians. For instance, several studies investigated GPT-3.5’s and GPT-4’s capability to generate discharge notes in tandem with various prompting strategies. In a UK pilot feasibility study, it was observed that a set of 25 AI-generated summaries were all deemed acceptable by general practitioners, compared to 23/25 (92%) of summaries written by junior doctors (Clough et al., 2024). Other studies similarly concluded that these proprietary models exhibit great potential and are able to generate acceptable discharge summaries with minimal misinformation (Kim et al., 2024; Waisberg et al., 2023). However, despite being able to increase efficiency and reduce the time required for documentation as compared to writing or dictating notes, instances of hallucination or omission of clinically significant facts were observed for certain discharge summaries involving complex surgeries. As such, the factual correctness of these large language models (LLMs) for specific generation tasks could be improved (Williams et al., 2024; Dubinski et al., 2024).

Based on this, some studies have focused on generating a particular section common to most discharge summaries – BHC – optimizing for correctness and faithfulness. The BHC is a succinct summary of a patient’s entire journey through the hospital and are embedded within complex discharge summaries. Efforts in compiling large-scale datasets for the generation of these BHC sections (Adams et al., 2021), including those with synthetic data (Adams et al., 2022), have led to subsequent contrastive learning methods for aligning generation models (Adams et al., 2023). Finally, methods leveraging heuristics to increase factuality (e.g., retrieval and ontology referencing) have also been developed (Adams et al., 2024; Hartman et al., 2023).

Some research has similarly centered on the Discharge Instructions section, sometimes known as the Patient Instructions section. This section is patient-facing and details instructions for the patient to continue their care at home, such as information on diet, therapies, and medications, as well as any details for follow-up appointments. Patient readability of this section is critical, and LLMs could be used to reformulate them into a more patient-friendly language (Zaretsky et al., 2024). Similar to the BHC, previous work also explored frameworks for the generation of faithful Patient Instructions (Liu et al., 2022).

3 RRG24: Radiology Report Generation

RRG24 was hosted on ViLMedic (Delbrouck et al., 2022b), a modular framework for vision-language multimodal research in medicine. The library contains reference implementations for state-of-the-art vision-language architectures for medicine and also hosts shared challenges in AI. A total of 201 submissions were received from across 8 teams.

3.1 Data

We curated Interpret-CXR, a large-scale collection of RRG datasets from the following five sources: MIMIC-CXR (Johnson et al., 2019), CheXpert (Irvin et al., 2019), PadChest (Bustos et al., 2020), BIMCV-COVID19 (Vayá et al., 2020), and OpenI (Demner-Fushman et al., 2016). The breakdown of Interpret-CXR, including details of the four splits used in RRG24 (Training, Validation, Public Test, and Hidden Test) are reported in Table 1.

3.2 Evaluation

We applied two types of metrics to evaluate different systems: n -gram-based and model-based metrics. For the former, we adopted BLEU-4 (Papineni et al., 2002) and ROUGE-L (Lin, 2004), whereas for the latter, we adopted BERTScore (Zhang et al., 2019), F1-CheXbert (Smit et al., 2020), and F1-RadGraph (Delbrouck et al., 2022a). To standardize the evaluation process, we used the same script from ViLMedic to evaluate all systems. By doing so, we avoid different teams using different versions or hyperparameters for a given metric – for example, some existing studies use differing versions of BERTScore, leading to inconsistent score reporting.

3.3 Results

The automatic results for the Findings and Impression sections are shown in Tables 2 and 3, respectively (Note: *iHealth-Chile-1* did not submit scores for Impression generation, and thus is not included in Table 3). We congratulate *e-Health CSIRO*, *MAIRA*, and *AIRI* for their outstanding performance on both Findings and Impression generation. It is also worth highlighting that the other teams (*Gla-AI4BioMed*, *SICAR*, *CID*, *iHealth-Chile-3&2*, and *iHealth-Chile-1*) designed novel solutions as well, providing insights for future research in this field beyond the competition. We also ran an evaluation using GREEN for the top 2 best-scoring systems (*e-Health CSIRO* and *MAIRA*) and recorded scores of 36.9 and 35.2, respectively, aligning with the leaderboard rankings¹.

3.4 Descriptions of Systems

3.4.1 e-Health CSIRO

e-Health CSIRO (Nicolson et al., 2024) integrated entropy regularization into self-critical sequence training to help maintain a higher entropy in the token distribution, preventing overfitting to common phrases and ensuring a broader exploration of the vocabulary during training. They applied this to a multimodal language model with RadGraph as the reward. Additionally, their model incorporated several other features: (i) the use of type embeddings to differentiate between Findings and Impression section tokens; and (ii) the use of a

¹We adopted GREEN instead of the naive GPT-4 pairwise comparison since Ostmeier et al. (2024) found GPT-4 to have low correlation with expert preference.

non-causal attention mask for image embeddings and a causal mask for report token embeddings.

3.4.2 MAIRA

MAIRA (Srivastav et al., 2024) combined a CXR-specific image encoder with a pre-trained LLM (Vicuna-7B-v1.5) via a multi-layer perceptron (MLP) adapter of 4 layers. The image encoder is a ViT-B model that leverages DINOv2, a state-of-the-art self-supervised learning method. Both the LLM and the adapter are fine-tuned in a single stage training setup for RRG. Their results indicated that joint training for Findings and Impression prediction improves the metrics for Findings generation. Additionally, incorporating lateral images alongside frontal images further enhances all metrics. They showed that scaling the model size from Vicuna-7B to Vicuna-13B also improves metrics. To handle multiple predictions for a study (as each study can have multiple frontal and/or lateral images), they utilized GPT-4 to select the best report.

3.4.3 AIRI

AIRI (Samokhin et al., 2024) utilized the LLaVA framework, where the vision encoder is a DINOv2 trained on medical data and the language decoder is a specialized biomedical LLM. They used the same model to generate both Impressions and Findings with different prompts: “Write findings for this X-ray.” or “Write impression for this X-ray.”. The system prompt from LLaVA-Med (Li et al., 2024) was also used.

3.4.4 Gla-AI4BioMed

Gla-AI4BioMed (Zhang et al., 2024) leveraged the Vicuna-7B architecture and integrated a CLIP image encoder with a fine-tuned LLM. The model underwent a two-stage training process, whereby chest X-ray features are initially aligned with the language model, and said model is subsequently fine-tuned for report generation. The model processed multiple images simultaneously by stitching them together, mimicking the workflow of radiology professionals.

3.4.5 SICAR

SICAR (Udomlaksakul et al., 2024) incorporated the SigLIP vision encoder and the Phi-2-2.7B language model to train an efficient RRG model. They also implemented a novel two-stage post-processing pipeline. They first enhanced the readability and clarity of the reports, then cross-verified the model outputs by integrating X-Raydar, an advanced X-ray classification model, addressing false negatives.

3.4.6 CID

CID (Liao et al., 2024) proposed a novel paradigm for incorporating graph structural data into the RRG model. Their approach involved predicting graph labels based on visual features and subsequently initiating the decoding process through a template injection conditioned on the predicted labels. These results provided preliminary

Table 1: Dataset Breakdown of Interpret-CXR for RRG24

Dataset	Training		Validation		Public Test		Hidden Test	
	Findings	Impression	Findings	Impression	Findings	Impression	Findings	Impression
PadChest	101,752	-	1,112	4,589	-	-	-	-
BIMCV-COVID19	45,525	-	1,202	-	-	-	-	-
CheXpert	45,491	181,619	2,641	-	-	-	-	-
OpenI	3,252	3,628	85	92	-	-	-	-
MIMIC-CXR	148,374	181,166	3,799	4,650	-	-	-	-
Total	344,394	366,413	8,839	9,331	2,692	2,967	1,063	1,428

Table 2: RRG24 Leaderboard for the Findings Section

Rank	Team	Overall Score \uparrow	Automatic Evaluation Metrics \uparrow				
			BLEU-4	ROUGE-L	BERTScore	F1-CheXbert	F1-RadGraph
1	e-Health CSIRO	35.56	11.68	26.16	53.80	57.49	28.67
2	MAIRA	35.08	11.24	26.58	54.22	57.87	25.48
3	AIRI	33.55	9.97	25.82	52.42	54.25	25.29
4	Gla-AI4BioMed	31.01	7.65	24.35	52.69	46.21	24.13
5	SICAR	30.93	6.62	23.66	50.74	49.00	24.62
6	CID	30.71	7.46	23.30	50.89	50.47	21.45
7	iHealth-Chile-3&2	23.38	4.81	15.96	44.03	33.69	18.41
8	iHealth-Chile-1	20.83	6.46	20.51	49.23	9.35	18.59

evidence for the feasibility of this new approach, which warrants further exploration in the future.

3.5 iHealth-Chile-3&2

iHealth-Chile-3&2 (Loch et al., 2024) focused on exploring various template-based strategies using predictions from multi-label image classifiers as input, which was inspired by prior work on template-based report generation. Two approaches were explored: (i) a straightforward implementation from Pino et al. (2021) directly; and (ii) replacing the fully connected layer with an attention-based pooling mechanism conditioned on a fact embedding.

3.6 iHealth-Chile-1

iHealth-Chile-1 (Campanini et al., 2024) developed a new strategy for in-context learning. Their system is formed using a vision-encoder, a vision-language connector or adapter, and a LLM able to process text and visual embeddings. They also designed an enriched prompt by combining a standard instruction (“Write the finding section of a chest x-ray radiology report”) with reports generated by a multi-label classifier and a group of template sentences.

3.7 Limitations & Challenges

The evaluation for medical text generation is challenging due to its domain-specific characteristics, making it difficult to measure performance as it relates to clinical utility. This challenge leveraged common metrics that are used by existing RRG studies. Unfortunately, these evaluations may be limited when considering the real-world clinical impact of the submitted systems.

4 “Discharge Me!”: Discharge Summary Generation

“Discharge Me!” was hosted on Codabench (Xu et al., 2022), an open source platform used to organize various tasks and benchmarks. A total of 211 submissions was received from across 16 teams.

4.1 Data

Participants were provided a dataset derived from the MIMIC-IV-Note module (Johnson et al., 2023). The modified and filtered dataset included 109,168 hospital admissions from the Emergency Department (ED), split into four sets (Training, Validation, Phase I Test, and Phase II Test) (Xu, 2024). Each visit includes chief complaints and diagnosis codes (either ICD-9 or ICD-10) documented by the ED², at least one radiology report, and a discharge summary with both BHC and Discharge Instructions sections.

The generation targets for the BHC were extracted from the full discharge notes using a complex regular expression strategy that searched for relevant section headers and new-line formatting characters. A similar strategy was used for Discharge Instructions; however, given that this section is usually located at the end of a discharge note as its very last section, extraction was more trivial. Samples where the extracted length of either section was shorter than 10 words were removed

²We assume ED diagnosis codes are available to the discharging clinician as ED documentation is likely to be complete at the time of discharge in most cases. However, we acknowledge that ICD codes may not necessarily be finalized, so they will be removed in future iterations of the shared task.

Table 3: RRG24 Leaderboard for the Impression Section

Rank	Team	Overall Score \uparrow	Automatic Evaluation Metrics \uparrow				
			BLEU-4	ROUGE-L	BERTScore	F1-CheXbert	F1-RadGraph
1	e-Health CSIRO	35.28	12.33	28.32	50.94	56.97	27.83
2	MAIRA	34.06	11.66	28.48	51.62	53.27	25.26
3	AIRI	32.98	10.91	27.46	49.55	52.32	24.67
4	SICAR	30.73	8.03	24.29	47.15	52.73	21.46
5	Gla-AI4BioMed	30.46	9.60	25.27	48.60	46.74	22.10
6	CID	25.21	7.13	20.41	43.67	39.64	15.19
7	iHealth-Chile-3&2	17.30	1.66	10.21	37.21	25.82	11.58

and deemed invalid. The complete breakdown of the dataset is available in Table 4.

Participants were allowed to incorporate external datasets, either publicly available or proprietary, as well as link additional patient data from other MIMIC-IV modules. Additionally, with the exception of the test dataset, participants were given the flexibility of using all or part of the provided dataset in any combination as they see fit.

4.2 Evaluation

4.2.1 Automatic Scoring

Automatic scoring took place on Codabench with a Python 3.9 environment. A hidden subset of 250 samples from the test datasets of the respective phases was used to evaluate the submissions. The metrics for this task were based on a combination of textual similarity (n -gram-based lexical metrics) and factual correctness of the generated text. Specifically, we considered the following metrics to automatically score submissions: BLEU-4 (Papineni et al., 2002), ROUGE-1/-2/-L (Lin, 2004), BERTScore (Zhang et al., 2019), METEOR (Banerjee and Lavie, 2005), AlignScore (Zha et al., 2023), and MEDCON (Van Veen et al., 2024).

Initially, submissions were scored on both target sections separately (BHC & Discharge Instructions). The mean across all test samples were computed for each metric, resulting in several performance scores for each of the two target sections (not reported on the leaderboard). Then, for each metric, we took the mean of the scores for each of the two target sections (reported under the metric name on the leaderboard). Finally, we computed the mean once again over all the metrics to arrive at a final overall system score (reported as Overall Score on the leaderboard).

For instance, given N samples, suppose s is defined as the score for a given sample for a given metric, then the mean across all samples for a particular target section, S , would be calculated by:

$$S = \sum_1^N (s_i) / N \quad (1)$$

We then calculated β , the mean of a given metric over both target sections, for each of the 8 metrics using:

$$\beta = (S_{BHC} + S_{DischargeInstructions}) / 2 \quad (2)$$

Finally, the overall system score was calculated by taking the mean of the 8 β values:

$$Overall = \sum_1^8 (\beta_i) / 8 \quad (3)$$

4.2.2 Clinician Scoring

At the end of the competition, the submissions from the top 6 best-scoring teams were reviewed by a group of six clinicians with diverse experiences in a broad range of specialties (two adult hospitalists, two clinical informatics fellows trained in pediatrics, a neurosurgeon, and a radiologist). Generated sections were evaluated for their completeness, correctness, and readability, as well as in a holistic comparison against the reference target sections (ground truth). In particular, completeness evaluates whether the generated text captures the clinically important information available in the reference text. In cases where there is inaccurate information, correctness specifies whether and how likely this mistake would lead to unintended impacts in future care. Readability was only evaluated by the clinicians for the BHC section as the intended audience of the Discharge Instructions section is the patient. Finally, the holistic comparison aimed to capture overall clinician preference.

Clinicians were presented with the reference target sections and the generated target sections side-by-side on a web-based survey dashboard hosted via Streamlit. Additionally, the full discharge summary was available in case reviewers required further context. They were then presented with a series of multiple-choice questions capturing each of the above criteria in a scale from 1 to 5, where 1 was the most negative option, and 5 was the most positive option.

Each clinician was provided with generated samples from three teams for evaluation. To minimize recall bias, we presented the generated submissions from all three teams consecutively in a randomized order for one particular sample, before moving onto the next.

Each team’s submission was evaluated by three separate clinician reviewers. Scores were averaged and several agreement and reliability scores were calculated, including Cohen’s Kappa and Fleiss Kappa for inter-observer agreement (McHugh, 2012; Landis and Koch, 1977), as well as the intraclass correlation coefficient (ICC) (Liljequist et al., 2019).

Table 4: Dataset Breakdown for “Discharge Me!”

Item	Total Count	Training	Validation	Phase I Test	Phase II Test
Hospital Visits	109,168	68,785	14,719	14,702	10,962
Discharge Summaries	109,168	68,785	14,719	14,702	10,962
Radiology Reports	409,359	259,304	54,650	54,797	40,608
ED Stays & Chief Complaints	109,403	68,936	14,751	14,731	10,985
ED Diagnoses	218,376	138,112	29,086	29,414	21,764

4.3 Results

4.3.1 Automatic Evaluation

Automatic scoring of the submissions took place on Codabench’s platform using queues connected to independent compute workers hosted on GCP. The final leaderboard on the Phase II Test set is available in Table 5.

A baseline performance was available for participants to benchmark their submissions. The baseline outputs were generated by a LLaMA-2-7B model finetuned on radiology reports from MIMIC-III (Johnson et al., 2016). While the system exhibited some clinical domain knowledge, it struggled due to the diverse formatting of discharge summaries, which greatly differed from that of the radiology reports in the training set. All submissions exceeded the baseline performance.

4.3.2 Clinician Evaluation

Overall clinician review scores are available in Table 6, and the specific rankings for the BHC and Discharge Instructions sections are shown in Tables 7 and 8, respectively (mean clinician scores are provided, along with their constituent scores in brackets). Interestingly, the rankings for the overall clinician review exactly reflected that of the automatic evaluation using the reported metrics.

Figure 4.3.2 illustrates the interobserver agreement between pairwise clinicians based on the Cohen’s Kappa statistic calculated for common submissions reviewed. As not all clinicians reviewed the same subset of submissions, a statistic could not be calculated for all reviewers (*i.e.*, reviewer #5 and #6 did not have any submissions in common). There was rather poor agreement between most clinicians, likely due to subjective aspects of the evaluation and varying clinician preference during the holistic comparison.

However, the Fleiss Kappa value indicated that the reviews for the top 6 best-scoring submissions, where each submission was reviewed by 3 individual clinicians, exhibited substantial to almost perfect agreement (Table 6). Moderate reliability was also observed for the review methodology, as inferred from the presented range of ICC values.

4.3.3 Readability of Discharge Instructions

As the Discharge Instructions section is intended for patients who many not have medical training and knowledge of clinical acronyms, we decided to skip the clinician review and opted for an evaluation using com-




Figure 1: Correlation heatmap visualizing interobserver agreement between clinician reviews. Cohen’s Kappa scores were computed between pairwise clinicians based on the respective common submission(s) reviewed.

mon readability scores: the Flesch Reading Ease score and the Flesch–Kincaid Grade Level (Friedman and Hoffman-Goetz, 2006).

The writing of patient-targeting notes at an appropriate readability level is crucial as it directly relates to patient comprehension, engagement, and adherence to treatment plans post-discharge. Several healthcare institutes have placed recommendations on the readability of patient-facing material. Specifically, the National Institutes of Health (NIH) and American Medical Association (AMA) encourage a reading grade level of not higher than sixth-grade, while the Centers for Disease Control and Prevention (CDC) suggests a reading grade level of lower than eighth-grade (Johnston et al., 2018; Cotugna et al., 2005; McCray, 2005; Burns et al., 2022).

A summary of the average readability metrics for the generated Discharge Instructions section is shown in Table 8. The readability of most submissions hovered around a reading grade level of seventh-grade, with the exception of one team at around the ninth-grade. The reference sections had a Flesch Reading Ease score of 61.81 (\pm 11.92) and a Flesch–Kincaid Grade Level of 8.16 (\pm 2.12). As such, all evaluated systems were able to reasonably re-create the readability of the reference sections, with several able to generate Discharge Instructions that are more understandable and in-line with established guidelines.

Table 5: “Discharge Me!” Automatic Scoring Leaderboard

Rank	Team	Overall Score \uparrow	Automatic Evaluation Metrics \uparrow							
			BLEU-4	ROUGE-1	ROUGE-2	ROUGE-L	BERTScore	METEOR	AlignScore	MEDCON
1	WisPerMed	0.332	0.124	0.453	0.201	0.308	0.438	0.403	0.315	0.411
2	HarmonAI Lab at Yale	0.300	0.106	0.423	0.180	0.284	0.412	0.381	0.265	0.353
3	aehrc	0.297	0.097	0.414	0.192	0.284	0.383	0.398	0.274	0.332
4	EPFL-MAKE	0.289	0.098	0.444	0.155	0.262	0.399	0.336	0.255	0.360
5	UF-HOBI	0.286	0.102	0.401	0.174	0.275	0.395	0.289	0.296	0.355
6	de ehren	0.284	0.097	0.404	0.166	0.265	0.389	0.376	0.231	0.339
7	DCT_PI	0.277	0.092	0.401	0.158	0.256	0.378	0.363	0.247	0.320
8	IgnitionInnovators	0.253	0.068	0.370	0.131	0.245	0.360	0.314	0.215	0.324
9	Shimo Lab	0.248	0.063	0.394	0.131	0.252	0.351	0.312	0.210	0.276
10	qub-cirdan	0.221	0.024	0.377	0.106	0.205	0.300	0.332	0.174	0.254
11	Roux-lette	0.206	0.030	0.319	0.084	0.182	0.289	0.287	0.195	0.265
12	UoG Siephers	0.191	0.017	0.341	0.109	0.209	0.268	0.247	0.143	0.193
13	mike-team	0.188	0.022	0.290	0.076	0.163	0.258	0.294	0.182	0.223
14	Ixa-UPV	0.183	0.016	0.259	0.057	0.144	0.282	0.284	0.210	0.215
15	MLBBIKABR	0.170	0.039	0.210	0.092	0.131	0.186	0.306	0.205	0.191
16	cyq	0.104	0.002	0.197	0.016	0.106	0.179	0.106	0.132	0.091

Table 6: “Discharge Me!” Clinician Scoring Leaderboard

Rank	Team	Average \uparrow	Fleiss Kappa	Intraclass Corr.
1	WisPerMed	3.375	0.781	0.336
2	HarmonAI Lab at Yale	2.903	0.944	0.656
3	aehrc	2.785	0.904	0.685
4	EPFL-MAKE	2.720	0.896	0.563
5	UF-HOBI	2.579	0.923	0.574
6	de ehren	2.335	0.908	0.740

4.4 Descriptions of Top Systems

A total of 12 system papers were received (Damm et al., 2024; Socrates et al., 2024; Wu et al., 2024; Lyu et al., 2024; He et al., 2024; Koontz et al., 2024; Guo et al., 2024; Liu et al., 2024; Frayling et al., 2024; Wendelken et al., 2024; Tang et al., 2024; Naskar et al., 2024). The top 6 best-scoring systems are detailed in this subsection.

4.4.1 WisPerMed

WisPerMed (Damm et al., 2024) investigated Dynamic Expert Selection (DES) consisting of a collection of LLMs fine-tuned and prompted for the task. They demonstrated that a DES system that chooses texts based on a specific length criteria performed the best on the given dataset. Thus, their objective with this strategy was to initially rank LLMs based on their archived overall scores. Subsequently, for each discharge summary, the generated sections (BHC & Discharge Instructions) from the best model that had a word count within the range of 100 to 180 words was selected. If no model generated a block of text with a word count within this range, the text with the minimum word count greater than 70 words was selected. In cases where no piece of text met these criteria (*i.e.*, shorter than 70 words), the text from the highest-ranked model was chosen. This approach emerged from the finding that longer pieces of medical text often led to hallucinations or repetitiveness.

4.4.2 HarmonAI Lab at Yale

The pipeline for **HarmonAI Lab at Yale** (Socrates et al., 2024) consisted of two BioBART-Large models.

The one generating BHC sections was trained on all the preceding text prior to the BHC, while the Discharge Instructions model was trained on the BHC. The BHC model had an increased training dataset size due to shuffling and recombining the provided datasets. Default hyperparameter settings were largely used for training, with the exception of a lower learning rate. Models were trained for 2 epochs. For generation, a 4-beam search and limited repeats with an n -gram size of 3 was employed. The minimum output length was set to 200 tokens based on the word count summary statistics and, and the maximum output token length was restricted to 1024 tokens due to the model specifications.

4.4.3 aehrc

aehrc (Liu et al., 2024) used the content in the discharge summary note prior to the target sections as input context for both training and inference. To better handle the distinctions between the two sections, the team trained two separate models to generate the BHC and the Discharge Instructions. Their best model was based on PRIMERA, which is an encoder-decoder language model that is capable of handling extended input contexts and generating longer outputs. This model offered a slight edge over fine-tuning popular decoder-based LLMs at the 7/8B parameter-level with LoRA, and was also significantly faster at inference. Beam search with a size of 4 was used for decoding.

4.4.4 EPFL-MAKE

EPFL-MAKE (Wu et al., 2024) mainly focused on the full-text available in the dataset as they believed that most of the useful information is hidden within. The text was used as an input into their system, which first extracted all sections that contained clinically useful information. The system then combined them into a new input. Some sections may have been removed if the new input was deemed too lengthy. The pre-processed input was then put into the medical LLM Meditron-7B, which is currently one of the top open-source medically pre-trained LLMs at the 7B level, to generate the BHC and Discharge Instructions sections.

Table 7: “Discharge Me!” Rankings based on Clinician Scoring of the Brief Hospital Course Section

Rank	Team	Average \uparrow	Clinician Evaluation Criteria \uparrow			
			Completeness	Correctness	Readability	Holistic Comparison
1	WisPerMed	3.29	3.67 (4.08 3.16 3.76)	3.67 (4.20 3.40 3.40)	3.37 (3.76 3.40 2.96)	2.44 (2.96 2.60 1.76)
2	EPFL-MAKE	2.58	3.29 (3.28 3.20 3.40)	2.83 (2.80 2.96 2.72)	2.53 (2.88 2.56 2.16)	1.65 (2.12 1.52 1.32)
3	UF-HOBI	2.49	2.48 (2.52 2.48 2.44)	3.36 (3.48 3.28 3.32)	2.71 (3.20 2.96 1.96)	1.41 (1.96 1.20 1.08)
4	HarmonAI Lab at Yale	2.44	3.52 (3.32 3.64 3.60)	2.59 (2.68 3.00 2.08)	2.11 (2.36 2.00 1.96)	1.53 (1.60 1.84 1.16)
5	de ehren	2.27	2.28 (2.36 2.32 2.16)	2.99 (3.12 3.24 2.60)	2.68 (2.72 2.84 2.48)	1.12 (1.16 1.20 1.00)
6	aehrc	2.10	2.31 (2.24 2.52 2.16)	3.05 (3.32 3.40 2.44)	1.96 (2.16 1.80 1.92)	1.09 (1.08 1.20 1.00)

Table 8: “Discharge Me!” Rankings based on Clinician Scoring of the Discharge Instructions Section

Rank	Team	Average \uparrow	Clinician Evaluation Criteria \uparrow			Flesch	Flesch-Kincaid
			Completeness	Correctness	Holistic Comparison	Reading Ease	Grade Level
1	aehrc	3.69	3.91 (3.80 4.40 3.52)	4.55 (4.52 4.48 4.64)	2.63 (2.48 3.24 2.16)	62.05 (\pm 10.04)	7.80 (\pm 1.76)
2	HarmonAI Lab at Yale	3.52	4.27 (3.88 4.40 4.52)	3.95 (3.84 3.88 4.12)	2.36 (2.36 2.40 2.32)	61.14 (\pm 14.52)	8.60 (\pm 4.19)
3	WisPerMed	3.49	3.95 (4.36 3.36 4.12)	4.00 (4.36 3.60 4.04)	2.53 (2.48 2.76 2.36)	63.35 (\pm 8.827)	7.48 (\pm 1.53)
4	EPFL-MAKE	2.91	3.45 (3.28 3.36 3.72)	3.41 (3.36 3.20 3.68)	1.87 (2.20 1.64 1.76)	58.72 (\pm 10.67)	9.04 (\pm 1.81)
5	UF-HOBI	2.70	3.01 (2.60 3.24 3.20)	3.29 (3.36 3.28 3.24)	1.79 (2.00 1.84 1.52)	66.73 (\pm 10.23)	6.96 (\pm 1.57)
6	de ehren	2.43	2.81 (2.84 3.12 2.48)	3.05 (3.36 3.12 2.68)	1.41 (1.44 1.60 1.20)	65.76 (\pm 8.706)	7.28 (\pm 1.84)

4.4.5 UF-HOBI

In their system, **UF-HOBI** (Lyu et al., 2024) employed two clinical LLMs that they have developed in their previous works, including an encoder-based model GatorTron (Yang et al., 2022) and a decoder-based model GatorTronGPT (Peng et al., 2023). The team adopted GatorTron to extract clinical concepts from the discharge summary notes, and utilized GatorTronGPT to generate the BHC and Discharge Instructions sections. GatorTron, which was fine-tuned on the 2010 i2b2 Challenge Named Entity Recognition (NER) dataset, was used to extract three categories of concepts (“TEST”, “PROBLEM”, and “TREATMENT”) from the discharge summary and radiology reports for each visit. The extracted concepts were then used to form the generation model input. Two GatorTronGPT models were then trained using the P-tuning strategy for the generation of the two respective target sections. The model inputs were thus the concepts extracted from the various other sections.

4.4.6 de ehren

de ehren utilized Meerkat-7B-v1.0, a compact, instruction-tuned medical AI system renowned for its advanced medical reasoning capabilities. Meerkat excelled in various medical Question Answering (QA) benchmarks, notably achieving a score of 74.3 on MedQA. To further scrutinize its performance in long-form text generation and summarization tasks within the clinical domain, the team selectively extracted key sections from discharge summaries to fine-tune the model with regards to the model’s attention window size.

4.5 Limitations & Challenges

A primary concern was the risk of data leakage due to the release of the test sets with ground truth sections. To

mitigate this, two test sets were released in two phases (one released at the start and one released much closer to the submission deadline), and the final evaluation was conducted on a hidden subset of 250 samples selected from the test datasets of the respective phases. This approach aimed to discourage participants from using the ground truth for model inference, or from optimizing systems for the tasks metrics throughout the entire duration of the competition. However, this method ultimately relies on the adherence of the participants to task guidelines.

The task also faced the challenge of dealing with inconsistently formatted free-text where ground truth generation targets are embedded within. The nature of clinical free-text can vary greatly, making it difficult to standardize inputs.

Furthermore, certain sections of the discharge summary appearing after the generation targets may not be reasonably available to the clinician at the time of discharge and the writing of the discharge summary. This presents a dilemma, as using such information would not accurately reflect the clinician’s workflow. Although teams were reminded to justify any decisions made regarding the use of discharge summary sections, it was challenging to moderate this aspect.

Another limitation was the need to select discharge summaries of a reasonable length to make clinician review feasible. This selection process may introduce a bias, as longer or more complex summaries that could benefit from automated generation might be excluded. There was also plausible comparison bias during clinician review as clinicians were asked to review submissions that could have varied greatly in quality. However, we aimed to reduce this by randomizing the order in which submissions were presented to the clinicians.

5 Conclusion

As seen from the scores of the participating models for both tasks, there is great complexity in generating coherent, accurate, and clinically relevant free-text reports. Several factors contribute to this, including the inherent variability and nuance of natural language used in clinical settings.

It may be worthwhile to consider alternative approaches for fully automated report generation, such as by pre-processing reports into structured formats prior to AI generation. By breaking down the report generation process into more manageable tasks, generation systems may be able to achieve higher accuracy and coherence in their outputs (Lederman et al., 2022). However, the standardization of formatting for these reports poses a significant challenge due to the diversity of writing styles and training among clinicians.

A previous study also explored the feasibility of generating hospital discharge summaries by tracing the source origin of medical expressions that make up the report (Ando et al., 2022). Interestingly, the analysis found that a significant portion of the discharge summary originates from external sources rather than inpatient records, such as past clinical records, referral notes, and the expertise of the writing clinician. This suggests that an end-to-end generation pipeline would depend on advanced data retrieval and may ultimately require some form of manual clinician oversight.

Ultimately, we hope that this challenge will bolster the efforts of the biomedical natural language processing community in developing effective solutions for clinical text generation. We believe this task could form a solid foundation for future work on generating entire radiology reports or discharge summaries, which would help significantly reduce the time clinicians spend on administrative tasks and improve patient care quality.

References

- Griffin Adams, Emily Alsentzer, Mert Ketenci, Jason Zucker, and Noémie Elhadad. 2021. [What’s in a Summary? Laying the Groundwork for Advances in Hospital-Course Summarization](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4794–4811, Online. Association for Computational Linguistics.
- Griffin Adams, Bichlien Nguyen, Jake Smith, Yingce Xia, Shufang Xie, Anna Ostropolets, Budhaditya Deb, Yuan-Jyue Chen, Tristan Naumann, and Noémie Elhadad. 2023. [What are the Desired Characteristics of Calibration Sets? Identifying Correlates on Long Form Scientific Summarization](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10520–10542, Toronto, Canada. Association for Computational Linguistics.
- Griffin Adams, Han-Chin Shing, Qing Sun, Christopher Winestock, Kathleen McKeown, and Noémie Elhadad. 2022. [Learning to Revise References for Faithful Summarization](#). ArXiv:2204.10290 [cs].
- Griffin Adams, Jason Zucker, and Noémie Elhadad. 2024. [SPEER: Sentence-Level Planning of Long Clinical Summaries via Embedded Entity Retrieval](#). ArXiv:2401.02369 [cs].
- Rana Alissa, Jennifer A. Hipp, and Kendall Webb. 2021. [Saving Time for Patient Care by Optimizing Physician Note Templates: A Pilot Study](#). *Frontiers in Digital Health*, 3:772356.
- Kenichiro Ando, Takashi Okumura, Mamoru Komachi, Hiromasa Horiguchi, and Yuji Matsumoto. 2022. [Is artificial intelligence capable of generating hospital discharge summaries from inpatient records?](#) *PLOS digital health*, 1(12):e0000158.
- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Shohei T. Burns, Nwamaka Amobi, Joshua Vic Chen, Meghan O’Brien, and Lawrence A. Haber. 2022. [Readability of Patient Discharge Instructions](#). *Journal of General Internal Medicine*, 37(7):1797–1798.
- Aurelia Bustos, Antonio Pertusa, Jose-Maria Salinas, and Maria De La Iglesia-Vaya. 2020. [Padchest: A large chest x-ray image dataset with multi-label annotated reports](#). *Medical image analysis*, 66:101797.
- Diego Campanini, Oscar Loch, Pablo Messina, Rafael Elberg, and Denis Parra. 2024. [ihealth-chile-1 at rrg24: In-context learning and finetuning of a large multimodal model for radiology report generation](#). In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- Zhihong Chen, Yan Song, Tsung-Hui Chang, and Xiang Wan. 2020. [Generating radiology reports via memory-driven transformer](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1439–1449.
- Reece Alexander James Clough, William Anthony Sparkes, Oliver Thomas Clough, Joshua Thomas Sykes, Alexander Thomas Steventon, and Kate King. 2024. [Transforming healthcare documentation: harnessing the potential of AI to generate discharge summaries](#). *BJGP open*, 8(1):BJGPO.2023.0116.
- Nancy Cotugna, Connie E. Vickery, and Kara M. Carpenter-Haefele. 2005. [Evaluation of literacy level of patient education pages in health-related journals](#). *Journal of Community Health*, 30(3):213–219.

- Hendrik Damm, Tabea M. G. Pakull, Bahadır Eryilmaz, Helmut Becker, Ahmad Idrissi-Yaghir, Henning Schäfer, Sergej Schultenkämper, and Christoph M. Friedrich. 2024. Wispermed at “discharge me!”: Advancing text generation in healthcare with large language models, dynamic expert selection, and priming techniques on mimic-iv. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- Jean-Benoit Delbrouck, Pierre Chambon, Christian Bluethgen, Emily Tsai, Omar Almusa, and Curtis Langlotz. 2022a. Improving the factual correctness of radiology report generation with semantic rewards. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4348–4360.
- Jean-benoit Delbrouck, Khaled Saab, Maya Varma, Sabri Eyuboglu, Pierre Chambon, Jared Dunnmon, Juan Zambrano, Akshay Chaudhari, and Curtis Langlotz. 2022b. Vilmedic: a framework for research at the intersection of vision and language in medical ai. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 23–34.
- Dina Demner-Fushman, Marc D Kohli, Marc B Rosenman, Sonya E Shooshan, Laritza Rodriguez, Sameer Antani, George R Thoma, and Clement J McDonald. 2016. Preparing a collection of radiology examinations for distribution and retrieval. *Journal of the American Medical Informatics Association*, 23(2):304–310.
- Huy M. Do, Lillian G. Spear, Moozhan Nikpanah, S. Mojdeh Mirmomen, Laura B. Machado, Alexandra P. Toscano, Baris Turkbey, Mohammad Hadi Bagheri, James L. Gulley, and Les R. Folio. 2020. [Augmented Radiologist Workflow Improves Report Value and Saves Time: A Potential Model for Implementation of Artificial Intelligence](#). *Academic Radiology*, 27(1):96–105.
- Daniel Dubinski, Sae-Yeon Won, Svorad Trnovec, Bedjan Behmanesh, Peter Baumgarten, Nazife Dinc, Juergen Konzalla, Alvin Chan, Joshua D. Bernstock, Thomas M. Freiman, and Florian Gessler. 2024. [Leveraging artificial intelligence in neurosurgery-unveiling ChatGPT for neurosurgical discharge summaries and operative reports](#). *Acta Neurochirurgica*, 166(1):38.
- Erlend Frayling, Jake Lever, and Graham McDonald. 2024. Uog siephers at discharge me!: Exploring ways to process multi-part electronic health records for sequence to sequence generation. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- Daniela B. Friedman and Laurie Hoffman-Goetz. 2006. [A systematic review of readability and comprehension instruments used for print and web-based cancer information](#). *Health Education & Behavior: The Official Publication of the Society for Public Health Education*, 33(3):352–373.
- Rui Guo, Greg Farnan, Niall McLaughlin, and Barry Devereux. 2024. Qub-cirdan at “discharge me!”: Zero shot discharge letter generation by open-source llm. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- Ibrahim Ethem Hamamci, Sezgin Er, and Bjoern Menze. 2024. [CT2Rep: Automated Radiology Report Generation for 3D Medical Imaging](#). ArXiv:2403.06801 [cs, eess] version: 1.
- Vince C. Hartman, Sanika S. Bapat, Mark G. Weiner, Babak B. Navi, Evan T. Sholle, and Thomas R. Champion. 2023. [A method to automate the discharge summary hospital course for neurology patients](#). *Journal of the American Medical Informatics Association: JAMIA*, 30(12):1995–2003.
- Michael Haycock, Laura Stuttaford, Oliver Ruscombe-King, Zoe Barker, Kathryn Callaghan, and Timothy Davis. 2014. [Improving the percentage of electronic discharge summaries completed within 24 hours of discharge](#). *BMJ Open Quality*, 3(1):u205963.w2604. Publisher: BMJ Open Quality Section: BMJ Quality Improvement Programme.
- Yunzhen He, Hiroaki Yamagiwa, and Hidetoshi Shimodaira. 2024. Shimo lab at “discharge me!”: Discharge summarization by prompt-driven concatenation of electronic health record sections. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- Wenjun Hou, Kaishuai Xu, Yi Cheng, Wenjie Li, and Jiang Liu. 2023. Organ: Observation-guided radiology report generation via tree reasoning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8108–8122.
- Jaeyoung Huh, Hyun Jeong Park, and Jong Chul Ye. 2023. [Breast Ultrasound Report Generation using LangChain](#). ArXiv:2312.03013 [cs, eess].
- Stephanie L Hyland, Shruthi Bannur, Kenza Bouzid, Daniel C Castro, Mercy Ranjit, Anton Schwaighofer, Fernando Pérez-García, Valentina Salvatelli, Shaury Srivastav, Anja Thieme, et al. 2023. Maira-1: A specialised large multimodal model for radiology report generation. *arXiv preprint arXiv:2311.13668*.
- Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpankaya, et al. 2019. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 590–597.

- Baoyu Jing, Pengtao Xie, and Eric Xing. 2018. On the automatic generation of medical imaging reports. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2577–2586.
- A Johnson, T Pollard, S Horng, LA Celi, and R Mark. 2023. MIMIC-iv-note: Deidentified free-text clinical notes (version 2.2). *physionet*.
- Alistair E. W. Johnson, Tom J. Pollard, Lu Shen, Liwei H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. 2016. **MIMIC-III, a freely accessible critical care database**. *Scientific Data*, 3(1):160035. Publisher: Nature Publishing Group.
- Alistair EW Johnson, Tom J Pollard, Seth J Berkowitz, Nathaniel R Greenbaum, Matthew P Lungren, Chihying Deng, Roger G Mark, and Steven Horng. 2019. MIMIC-cxr, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific data*, 6(1):317.
- David Johnston, Owen McMurray, Michael McKee, Michael McConville, and Niall Leonard. 2018. **‘DISCHARGE LETTER QUALITY; HOW TO HELP BOTH JUNIOR DOCTORS AND GPS?’**. *The Ulster Medical Journal*, 87(2):130.
- Hanjae Kim, Hee Min Jin, Yoon Bin Jung, and Seng Chan You. 2024. **Patient-Friendly Discharge Summaries in Korea Based on ChatGPT: Software Development and Validation**. *Journal of Korean Medical Science*, 39(16):e148.
- Jordan Koontz, Maite Oronoz, and Alicia Pérez. 2024. Ixa-med at discharge me! retrieval-assisted generation for streamlining discharge documentation. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- J. R. Landis and G. G. Koch. 1977. The measurement of observer agreement for categorical data. *Biometrics*, 33(1):159–174.
- Asher Lederman, Reeva Lederman, and Karin Verspoor. 2022. **Tasks as needs: reframing the paradigm of clinical natural language processing research for real-world decision support**. *Journal of the American Medical Informatics Association*, 29(10):1810–1817.
- Chunyu Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. 2024. Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *Advances in Neural Information Processing Systems*, 36.
- Yuxiang Liao, Yuanbang Liang, Yipeng Qin, Hantao Liu, and Irena Spasić. 2024. Cid at rrg24: Attempting in a conditionally initiated decoding of radiology report generation with clinical entities. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- David Liljequist, Britt Elfving, and Kirsti Skavberg Roaldsen. 2019. **Intraclass correlation – A discussion and demonstration of basic features**. *PLoS ONE*, 14(7):e0219854.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Fenglin Liu, Xian Wu, Shen Ge, Wei Fan, and Yuexian Zou. 2021. Exploring and distilling posterior and prior knowledge for radiology report generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13753–13762.
- Fenglin Liu, Bang Yang, Chenyu You, Xian Wu, Shen Ge, Zhangdaihong Liu, Xu Sun, Yang Yang, and David A. Clifton. 2022. **Retrieve, Reason, and Refine: Generating Accurate and Faithful Patient Instructions**. In *Proceedings of the Neural Information Processing Systems*.
- Jinghui Liu, Aaron Nicolson, Jason Dowling, Bevan Koopman, and Anthony Nguyen. 2024. e-health csiro at “discharge me!” 2024: Generating discharge summary sections with fine-tuned language models. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- Oscar Loch, Pablo Messina, Rafael Elberg, Diego Campanini, Álvaro Soto, René Vidal, and Denis Parra. 2024. ihealth-chile-3&2 at rrg24: Template based report generation. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- Samira Loveymi, Mir Hossein Dezfoulian, and Muharram Mansoorizadeh. 2021. **Automatic Generation of Structured Radiology Reports for Volumetric Computed Tomography Images Using Question-Specific Deep Feature Extraction and Learning**. *Journal of Medical Signals and Sensors*, 11(3):194–207.
- Mengxian Lyu, Cheng Peng, Daniel Paredes, Ziyi Chen, Aokun Chen, Jiang Bian, and Yonghui Wu. 2024. Uf-hobi at “discharge me!”: A hybrid solution for discharge summary generation through prompt-based tuning of gatortrongpt models. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- Alexa T. McCray. 2005. **Promoting Health Literacy**. *Journal of the American Medical Informatics Association : JAMIA*, 12(2):152–163.
- Mary L. McHugh. 2012. **Interrater reliability: the kappa statistic**. *Biochemia Medica*, 22(3):276–282.

- Abir Naskar, Jane Hocking, Patty Chondros, Douglas Boyle, and Mike Conway. 2024. Mlbmikabr at "discharge me!": Concept based clinical text description generation. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- Aaron Nicolson, Jinghui Liu, Jason Dowling, Anthony Nguyen, and Bevan Koopman. 2024. e-health csiro at rrg24: Entropy-augmented self-critical sequence training for radiology report generation. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- Sophie Ostmeier, Justin Xu, Zhihong Chen, Maya Varma, Louis Blankemeier, Christian Bluethgen, Arne Edward Michalson, Michael Moseley, Curtis Langlotz, Akshay S. Chaudhari, and Jean-Benoit Delbrouck. 2024. [GREEN: Generative Radiology Report Evaluation and Error Notation](#). ArXiv:2405.03595 [cs].
- Ting Pang, Peigao Li, and Lijie Zhao. 2023. [A survey on automatic generation of medical imaging reports based on deep learning](#). *BioMedical Engineering OnLine*, 22(1):48.
- Kishore Papineni, Salim Roukos, Todd Ward, and Weijing 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.
- Cheng Peng, Xi Yang, Aokun Chen, Kaleb E. Smith, Nima PourNejatian, Anthony B. Costa, Cheryl Martin, Mona G. Flores, Ying Zhang, Tanja Magoc, Gloria Lipori, Duane A. Mitchell, Naykky S. Ospina, Mustafa M. Ahmed, William R. Hogan, Elizabeth A. Shenkman, Yi Guo, Jiang Bian, and Yonghui Wu. 2023. [A study of generative large language model for medical research and healthcare](#). *npj Digital Medicine*, 6(1):1–10. Publisher: Nature Publishing Group.
- Pablo Pino, Denis Parra, Cecilia Besa, and Claudio Lagos. 2021. Clinically correct report generation from chest x-rays using templates. In *Machine Learning in Medical Imaging: 12th International Workshop, MLMI 2021, Held in Conjunction with MICCAI 2021, Strasbourg, France, September 27, 2021, Proceedings 12*, pages 654–663. Springer.
- V. Samokhin, M. Munkhoeva, and D. Umerenkov. 2024. Airi at rrg24: Llava with specialised encoder and decoder. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- Hoo-Chang Shin, Kirk Roberts, Le Lu, Dina Demner-Fushman, Jianhua Yao, and Ronald M Summers. 2016. Learning to read chest x-rays: Recurrent neural cascade model for automated image annotation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2497–2506.
- Akshay Smit, Saahil Jain, Pranav Rajpurkar, Anuj Pareek, Andrew Y Ng, and Matthew Lungren. 2020. Combining automatic labelers and expert annotations for accurate radiology report labeling using bert. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1500–1519.
- Vimig Socrates, Thomas Huang, Xuguang Ai, Soraya Fereydooni, Qingyu Chen, R. Andrew Taylor, and David Chartash. 2024. Yale at "discharge me!": Evaluating constrained generation of discharge summaries with unstructured and structured information. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- Shaury Srivastav, Mercy Ranjit, Fernando Pérez-García, Kenza Bouzid, Shruthi Bannur, Daniel C. Castro, Anton Schwaighofer, Harshita Sharma, Maximilian Ilse, Valentina Salvatelli, Sam Bond-Taylor, Fabian Falck, Anja Thieme, Hannah Richardson, Matthew P. Lungren, Stephanie L. Hyland, and Javier Alvarez-Valle. 2024. Maira at rrg24: A specialised large multimodal model for radiology report generation. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- An Quang Tang, Xiuzhen Zhang, and Minh Ngoc Dinh. 2024. Ignitioninnovators at "discharge me!": Chain-of-thought instruction finetuning large language models for discharge summaries. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- Omkar Thawkar, Abdelrahman Shaker, Sahal Shaji Mullappilly, Hisham Cholakkal, Rao Muhammad Anwer, Salman Khan, Jorma Laaksonen, and Fahad Shahbaz Khan. 2023. Xraygpt: Chest radiographs summarization using medical vision-language models. *arXiv preprint arXiv:2306.07971*.
- Tao Tu, Shekoofeh Azizi, Danny Driess, Mike Schaekermann, Mohamed Amin, Pi-Chuan Chang, Andrew Carroll, Charles Lau, Ryutaro Tanno, Ira Ktena, et al. 2024. Towards generalist biomedical ai. *NEJM AI*, 1(3):AIoa2300138.
- Kiartnarin Udomlapsakul, Parinthapat Pengpun, Tossaporn Saengja, Kanyakorn Veerakanjana, Krittamate Tiankanon, Pitikorn Khlaisamniang, Pasit Supholkhan, Amrest Chinkamol, Pubordee Ausavavirojekul, Hirunkul Phimsiri, Tara Sripo, Chiraphat Boonnag, Trongtum Tongdee, Thanongchai Siriapisith, Pairash Saiviroonporn, Jiramet Kinchawat, and Piyalitt Ittichaiwong. 2024. Sicar at rrg2024: Gpu poor's guide to radiology report generation. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.

- Dave Van Veen, Cara Van Uden, Louis Blankemeier, Jean-Benoit Delbrouck, Asad Aali, Christian Bluethgen, Anuj Pareek, Malgorzata Polacin, 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*.
- Maria De La Iglesia Vayá, Jose Manuel Saborit, Joaquim Angel Montell, Antonio Pertusa, Aurelia Bustos, Miguel Cazorla, Joaquin Galant, Xavier Barber, Domingo Orozco-Beltrán, Francisco García-García, et al. 2020. Bimcv covid-19+: a large annotated dataset of rx and ct images from covid-19 patients. *arXiv preprint arXiv:2006.01174*.
- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. [CIDER: Consensus-based Image Description Evaluation](#). ArXiv:1411.5726 [cs].
- Ethan Waisberg, Joshua Ong, Mouayad Masalkhi, Sharif Amit Kamran, Nasif Zaman, Prithul Sarker, Andrew G. Lee, and Alireza Tavakkoli. 2023. [GPT-4: a new era of artificial intelligence in medicine](#). *Irish Journal of Medical Science*, 192(6):3197–3200.
- Katharine Weetman, Rachel Spencer, Jeremy Dale, Emma Scott, and Stephanie Schnurr. 2021. [What makes a “successful” or “unsuccessful” discharge letter? Hospital clinician and General Practitioner assessments of the quality of discharge letters](#). *BMC Health Services Research*, 21:349.
- S. Wendelken, A. Antony, R. Korutla, B. Pachipala, J. Shanahan, and W. Saba. 2024. Roux-lette at “discharge me!”: Reducing ehr chart burden with a simple, scalable, clinician-driven ai approach. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- Christopher Y. K. Williams, Jaskaran Bains, Tianyu Tang, Kishan Patel, Alexa N. Lucas, Fiona Chen, Brenda Y. Miao, Atul J. Butte, and Aaron E. Kornblith. 2024. [Evaluating Large Language Models for Drafting Emergency Department Discharge Summaries](#). *medRxiv: The Preprint Server for Health Sciences*, page 2024.04.03.24305088.
- Haotian Wu, Paul Boulenger, Antonin Faure, Berta Céspedes, Farouk Boukil, Nastasia Morel, Zeming Chen, and Antoine Bosselut. 2024. Epfl-make at “discharge me!”: An llm system for automatically generating discharge summaries of clinical electronic health record. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- Justin Xu. 2024. [Discharge Me: BioNLP ACL’24 Shared Task on Streamlining Discharge Documentation](#).
- Zhen Xu, Sergio Escalera, Adrien Pavão, Magali Richard, Wei-Wei Tu, Quanming Yao, Huan Zhao, and Isabelle Guyon. 2022. [Codabench: Flexible, easy-to-use, and reproducible meta-benchmark platform](#). *Patterns*, 3(7):100543.
- Shaokang Yang, Jianwei Niu, Jiyan Wu, Yong Wang, Xuefeng Liu, and Qingfeng Li. 2021. Automatic ultrasound image report generation with adaptive multimodal attention mechanism. *Neurocomputing*, 427:40–49.
- Xi Yang, Aokun Chen, Nima PourNejatian, Hoo Chang Shin, Kaleb E. Smith, Christopher Parisien, Colin Compas, Cheryl Martin, Anthony B. Costa, Mona G. Flores, Ying Zhang, Tanja Magoc, Christopher A. Harle, Gloria Lipori, Duane A. Mitchell, William R. Hogan, Elizabeth A. Shenkman, Jiang Bian, and Yonghui Wu. 2022. [A large language model for electronic health records](#). *npj Digital Medicine*, 5(1):1–9. Publisher: Nature Publishing Group.
- Jonah Zaretsky, Jeong Min Kim, Samuel Baskharoun, Yunan Zhao, Jonathan Austrian, Yindalon Aphinyanaphongs, Ravi Gupta, Saul B. Blecker, and Jonah Feldman. 2024. [Generative Artificial Intelligence to Transform Inpatient Discharge Summaries to Patient-Friendly Language and Format](#). *JAMA Network Open*, 7(3):e240357.
- Xianhua Zeng, Li Wen, Banggui Liu, and Xiaojun Qi. 2020. [Deep learning for ultrasound image caption generation based on object detection](#). *Neurocomputing*, 392:132–141.
- Yuheng Zha, Yichi Yang, Ruichen Li, and Zhiting Hu. 2023. [AlignScore: Evaluating factual consistency with a unified alignment function](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11328–11348, Toronto, Canada. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.
- Xi Zhang, Zaiqiao Meng, Jake Lever, and Edmond S.L. Ho. 2024. [Gla-ai4biomed at rrg24: Visual instruction-tuned adaptation for radiology report generation](#). In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- Yixiao Zhang, Xiaosong Wang, Ziyue Xu, Qihang Yu, Alan Yuille, and Daguang Xu. 2020. [When radiology report generation meets knowledge graph](#). In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 12910–12917.
- Zizhao Zhang, Yuanpu Xie, Fuyong Xing, Mason McGough, and Lin Yang. 2017. [Mdnnet: A semantically and visually interpretable medical image diagnosis network](#). In *Proceedings of the IEEE conference*

on computer vision and pattern recognition, pages
6428–6436.