



OPEN ACCESS

EDITED BY

Lei Bi,
Shanghai Jiao Tong University, China

REVIEWED BY

Geng Chen,
Northwestern Polytechnical University, China
Hongyu Wang,
Northwestern Polytechnical University, China

*CORRESPONDENCE

Nurbanu Aksoy
✉ nurbanudemiraksoy@gmail.com

†PRESENT ADDRESSES

Alejandro F. Frangi,
Christabel Pankhurst Institute, The University
of Manchester, Manchester, United Kingdom;
NIHR Manchester Biomedical Research
Centre, Manchester Academic Health Science
Centre, Manchester, United Kingdom

†These authors share last authorship

RECEIVED 16 November 2023

ACCEPTED 25 January 2024

PUBLISHED 15 February 2024

CITATION

Aksoy N, Sharoff S, Baser S, Ravikumar N and
Frangi AF (2024) Beyond images: an integrative
multi-modal approach to chest x-ray report
generation.
Front. Radiol. 4:1339612.
doi: 10.3389/fradi.2024.1339612

COPYRIGHT

© 2024 Aksoy, Sharoff, Baser, Ravikumar and
Frangi. This is an open-access article
distributed under the terms of the [Creative
Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The
use, distribution or reproduction in other
forums is permitted, provided the original
author(s) and the copyright owner(s) are
credited and that the original publication in
this journal is cited, in accordance with
accepted academic practice. No use,
distribution or reproduction is permitted
which does not comply with these terms.

Beyond images: an integrative multi-modal approach to chest x-ray report generation

Nurbanu Aksoy^{1*}, Serge Sharoff², Selcuk Baser³,
Nishant Ravikumar^{1†} and Alejandro F. Frangi^{4,5††}

¹Center for Computational Imaging & Simulation Technologies in Biomedicine, School of Computing, University of Leeds, Leeds, United Kingdom, ²School of Languages, University of Leeds, Leeds, United Kingdom, ³Kastamonu Training and Research Hospital, Kastamonu, Türkiye, ⁴Medical Imaging Research Centre, KU Leuven, Leuven, Belgium, ⁵Alan Turing Institute, London, United Kingdom

Image-to-text radiology report generation aims to automatically produce radiology reports that describe the findings in medical images. Most existing methods focus solely on the image data, disregarding the other patient information accessible to radiologists. In this paper, we present a novel multi-modal deep neural network framework for generating chest x-rays reports by integrating structured patient data, such as vital signs and symptoms, alongside unstructured clinical notes. We introduce a conditioned cross-multi-head attention module to fuse these heterogeneous data modalities, bridging the semantic gap between visual and textual data. Experiments demonstrate substantial improvements from using additional modalities compared to relying on images alone. Notably, our model achieves the highest reported performance on the ROUGE-L metric compared to relevant state-of-the-art models in the literature. Furthermore, we employed both human evaluation and clinical semantic similarity measurement alongside word-overlap metrics to improve the depth of quantitative analysis. A human evaluation, conducted by a board-certified radiologist, confirms the model's accuracy in identifying high-level findings, however, it also highlights that more improvement is needed to capture nuanced details and clinical context.

KEYWORDS

report generation, transformers, cross attention, multi-modal data, x-ray, deep learning

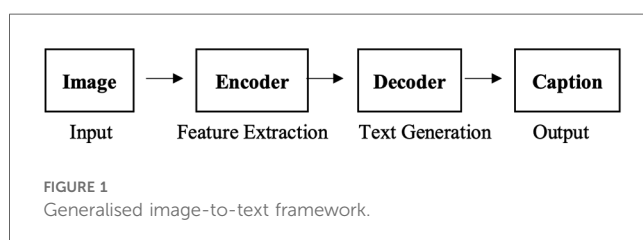
1 Introduction

The use of medical imaging is widespread across various branches of health sciences for the purpose of diagnosing diseases, developing effective treatment plans, providing patient care, and predicting disease outcomes. Radiologists are responsible for interpreting the medical images and creating a full-text radiology report that is based on their findings along with other relevant clinical data and information, such as patient demographics, symptoms, and pre-existing/existing medical conditions. These reports must be complete, accurate, and produced in a short amount of time while adhering to a specific format. In clinical settings, chest x-rays (CXR) are the most commonly used medical imaging techniques and are usually the first step in evaluating patients for various lung diseases. The reports generated from CXR examinations typically include the radiologists' observations categorised as "findings" and "impressions" and indicate normal and abnormal features in the images. Composing these detailed reports requires a significant amount of knowledge and

experience and can be time-consuming and prone to errors. By providing radiologists with a baseline analysis to validate and amend as needed, automation can reduce repetitive workflows. This would allow radiologists to focus their expertise on higher-level clinical thinking and quality assurance.

In the field of medical imaging informatics, previous studies have developed techniques to automate the generation of radiology reports (1, 2). The majority of current deep learning approaches use networks that feature a convolutional encoder and recurrent (3–5) or transformer decoder (6, 7), which were originally designed for the task of image captioning (Figure 1). Although these two tasks share similarities in terms of input and output modalities, there are some key differences. Radiology reports are in the form of detailed paragraphs rather than brief captions, and they must be comprehensive and include specific medical details. Additionally, interpreting medical images can be challenging due to subtle variations in the image and report and also generating a description for a medical image often necessitates supplementary information beyond what is visible in the image. For instance, in certain cases, while similarities in medical imaging between males and females are nearly identical in terms of visual patterns, differences in patient demographics have a noteworthy clinical impact on the assessment and diagnosis. However, current report generation methods for CXRs solely consider the radiology image as input and disregard the non-imaging information that radiologists have access to during image interpretation. Only a limited number of studies integrate additional data into the network such as medical concepts, high-level contexts or categories of the images/reports. While these methods have shown some level of success, they mainly focus on enhancing the model with data derived from existing semantics rather than supplementing the training context with additional data. Furthermore, as CXR images are 2D projections of 3D objects, important/relevant information is lost, leading to semantic gaps in the data available for learning by networks or algorithms. Therefore, we hypothesise that combining multiple data sources that provide different perspectives on the patient's condition, is beneficial for generating more informative and accurate reports (using data-driven learning-based approaches), compared to using CXR images alone.

The rest of the paper is organised as follows: Section 2 presents relevant existing studies and strategies, and Section 3 introduces the data and outlines the proposed methodology. Section 4 provides details on the implementation and presents both quantitative and qualitative results. Finally, Section 5 conducts a discussion and draws a series of conclusions regarding the study.



2 Relevant literature

2.1 Image captioning methodologies

Visual Captioning, also known as Image Captioning, is a popular task that involves generating descriptive natural language captions for images. It requires the integration of computer vision and natural language processing, drawing significant attention from the artificial intelligence community. Deep learning techniques, particularly encoder-decoder models using Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), have been commonly employed for this task (8, 9). However, these RNN-based models suffer from the issue of vanishing and exploding gradients due to limited access to previous inputs, affecting their performance (10). To address this limitation, recent research has shifted towards utilising transformer-based architectures (11), originally successful in the field of natural language processing (NLP). Transformers leverage self-attention mechanisms, allowing for better parallelisation and learning relationships between words in a sequence (12). Unlike RNNs, transformers do not rely on recurrence, enabling faster and more effective learning by including more context in the network.

The transformer-based encoder-decoder architecture that is most commonly used for image captioning consists of three main components: a model for extracting visual features, a transformer-based encoder, and a transformer-based decoder. To extract high-level features, pre-trained CNN models are usually employed. However, in this approach, the output of the visual model is passed through a transformer-based encoder to map the visual features and produce a sequence of image representations. The transformer-based decoder then takes in the encoder's output to generate a corresponding caption for the given image. One model, called Captioning Transformer (CT) (13), utilised a ResNeXt (14) CNN model as an encoder and a Transformer as a decoder. Another study (15) used a Transformer as the decoder along with the ResNet CNN model and improved the network with a combination of spatial and adaptive attention. Another study (16) enhanced the vanilla Transformer architecture with Entangled Attention (ETA) and Gated Bilateral Controller (GBC), enabling the processing of semantic and visual concepts concurrently. A Meshed-Memory Transformer model (17) was introduced that is a fully-attentive model including a Memory-Augmented Encoder enriched with learnable keys and values and a gating mechanism for mesh connectivity. The model also includes a Meshed Decoder that connects all encoding layers. A separate study proposed the Caption Transformer (18) (CPTR) model, which is a full Transformer network without any convolutional operation in the encoder. Unlike previous models, the CPTR model uses the raw image, divides it into N patches, reshapes the patches into vectors and learns features from them, with positional encoding, using the transformer encoder.

2.2 Chest-xray report generation

Recent approaches to automatic radiology/CXR image report generation are predominantly based on deep learning and typically adopt an image captioning approach, utilising a combination of convolution and recurrent neural networks. For example, Jing et al. (3), used a pre-trained VGG-19 model to learn visual features, which are then used to predict relevant tags for any given chest x-ray. These tags become semantic features in the network, which alongside visual features, are fed into a Co-Attention Network. Subsequently, a Hierarchical LSTM uses the context vector provided by the Co-Attention Network to generate a description for the given x-ray. While the model achieved promising results, there were concerns around the repetitive sentences in reports and the inconsistency in generating text for the same patient during inference, where the network produced different outcomes. Another study (4) addressed this issue by enhancing the pre-trained Resnet-152 encoder using multi-view content and incorporating a sentence decoder to generate a report. Multi-view approach mitigated the problem of report variability for the same patient. Additionally, they used the first predicted sentence as a joint input alongside image encoding, ensuring consistency in the results.

An alternative proposal (5) was posited, wherein they chose to pre-train their multi-view encoder from randomly initialised weights using the CheXpert dataset (19), rather than relying on pre-trained models from ImageNet. To improve the efficacy of the decoder, they extracted normalised medical concepts from radiology reports using Semrep.¹ These extracted concepts were embedded into the decoder in two ways: (1) they concatenated the concept embedding with the encoder output before feeding to the decoder, providing explicit semantic information, and (2) they used a concept-aware attention mechanism that attends over the embedded concepts when generating words, enabling the decoder to focus on relevant medical terms. By enriching the decoder with explicit medical concept knowledge, their model could generate reports with more accurate and meaningful terminology aligned with the clinical finding descriptions. Medical concepts were also employed in another study (1), where, the authors introduced a reinforcement learning-based reward for concept extraction to obtain more accurate and precise concepts. This reward encouraged the model to extract concepts that have a higher likelihood of being mentioned in the radiologist's report. Compared to previous work that extracted concepts without optimisation, their approach achieved higher precision and coverage of concepts that radiologists tend to use in real reports. However, the authors acknowledged their generated reports still lacked some descriptive informativeness compared to ground truth. Their analysis found that only extracting concepts from previous reports limits the diversity of expressions in generated reports.

A study (2) highlighted the frequent discrepancy between the formats of normal and abnormal radiology reports, with abnormal reports indicating the suspicion of an abnormality or pathology. To address this issue, they categorised reports as either normal or abnormal based on the content of the report text. For the report generation process, they adopted a two-stage approach. First, they generated the "Findings" section, which describes the visual observations made by radiologists during the examination of medical images. This was achieved by leveraging visual features extracted from a CNN model and relevant report texts. Subsequently, they summarised the generated "Findings" to produce the "Impression" section, which provides a summarised interpretation of the radiologist's observations. A key contribution of their study was showing that conditioning the text generation on the report type (normal or abnormal) improved the clinical validity and alignment with real radiology reporting practices.

Recent studies have also capitalised on Transformer models for medical report generation, after having achieved success in text generation based on non-linguistic representation. One such study (20) constructed a hierarchical Transformer model which features a novel encoder capable of extracting regions of interest from the original image via a region detector, and subsequently, utilising these regions to obtain visual representations. Additionally, another study (6) introduced a medical report generator utilising a memory-driven Transformer. They proposed a relational memory (RM) module to retain knowledge from previous cases, thereby enabling the generator model to remember similar reports when generating current reports. Another study (7) proposed a progressive Transformer-based framework for report generation, which generates high-level context from the given x-ray and then employs the Transformer architecture to convert this context into a radiology report. This model comprises a pre-trained CNN as a visual backbone, a mesh-memory Transformer (17) as a visual-language model, and BART (21) as a language model.

With increasing interest in this application domain, studies have become more attentive to the distinctions between image captioning and report generation tasks. As a result, researchers have begun to develop more knowledge-informed networks tailored specifically to the task of image-guided radiology report generation. The paper (22) introduces a task-aware framework that is designed to be adaptable to different imaging types and medical scenarios. It prioritises understanding specific diagnostic tasks related to various medical conditions, ensuring accurate and contextually relevant report generation. Another study (23) highlights the significance of both input-independent general medical knowledge and input-dependent specific contextual information in generating accurate chest radiology reports. They proposed a knowledge-enhanced method that leverages these information along with visual features to improve the quality and accuracy of generated reports for chest x-rays. Recently, another study (24) introduced a technique called multi-modal contrastive learning, which aims to enhance the synergy between different modalities of data. By leveraging contrastive learning, the proposed method aligns and embeds visual and textual representations in a shared space, facilitating the generation of more informative and accurate radiology reports.

¹<https://semrep.nlm.nih.gov/>

2.3 Fusion strategies

Data fusion refers to the integration of different data modalities that provide separate perspectives on a problem to be addressed, and using multiple modalities has the potential to decrease the number of errors compared to approaches that only use one type of data (25). Deep learning fusion strategies can be broadly classified into three categories: early fusion, late/decision fusion, and hybrid/joint fusion. In the process of early fusion, the original or transformed features are combined at the input level before being fed into a single model that can handle all the information. There are various methods of joining data, but early fusion commonly involves concatenation or pooling. In late fusion, the input data is processed independently through separate networks. The outputs from these networks are then combined at a later stage to form a joint decision. Late fusion strategies learn modality-specific features separately and then integrate them downstream in the model (e.g., just before the prediction/output layer). Lastly, joint fusion involves combining the features extracted from different modalities at different stages of the network architecture.

Within the medical imaging field, the utilisation of multi-modal data fusion approaches has the potential to enhance performance in addressing complex tasks that exceed the capabilities of a single imaging modality. Concentrating on chest x-ray modality, multiple tasks such as image classification, image retrieval, and modality translation have leveraged data fusion strategies. For example, one study (26) introduced a CNN-RNN architecture called the text-image-embedding network (TieNet) to extract discriminative representations of both chest radiographs and their accompanying reports by combining visual and textual information through joint fusion. The experimental results indicate that TieNet's multimodal approach outperforms its unimodal counterpart in multi-label disease classification. Another study (27) employed a semi-supervised approach to train the network on chest radiographs and associated radiology reports to evaluate the severity of pulmonary edema. This study demonstrated that joint learning of image-text representations enhances the performance of models designed to predict the severity of pulmonary edema, compared with supervised models that relied solely image-derived features. A paper (28) discussed the challenge of integrating data from different sources in healthcare due to asynchronous collection of modalities. They proposed an LSTM-based fusion module, called MedFuse, that accommodates uni-modal and multi-modal input for mortality prediction and phenotype classification tasks. In contrast with intricate multi-modal fusion techniques, MedFuse yields considerably better performance on the fully paired test set, furthermore, it demonstrates robustness when dealing with the partially paired test set, which includes instances of missing chest x-ray images.

This paper introduces a novel multi-modal deep neural network for generating radiology reports based on quantitative image analysis, patient demographic information, and other clinical data collected during a patient's stay. The main objective is to generate a consistent and comprehensive report that adheres to the format used in real-world clinical practice. Our contributions are as follows:

To the best of our knowledge, this is the first attempt to generate an automatic CXR report by integrating patient information and clinical data obtained from CXR exams, which is not typically included in radiology reports. We propose a novel conditioned cross-multi-head attention module to fuse structured data, unstructured text and visual information. Additionally, we employed human evaluation and clinical semantic similarity measurement [Bio-ClinicalBERT Score (1)] alongside word-overlap metrics to improve the depth of quantitative analysis. Our experiments demonstrate that the incorporation of additional data not explicitly stated in the report enhances the model's performance. Our proposed model achieves the best performance on the ROUGE-L metric (29) when compared to similar state-of-the-art studies.

3 Materials and methods

3.1 Data

The dataset used in this study was created by leveraging three openly accessible databases, namely MIMIC-CXR, MIMIC-IV, and MIMIC-IV-ED. MIMIC-CXR (version 2.0) encompasses a vast collection of 377,110 CXR images captured from multiple views, together with 227,835 de-identified radiology reports, pertaining to 63,473 patients. Each report contains several sections, such as "examination", "indication", "technique", "comparison", "findings", and "impressions". Meanwhile, MIMIC-IV (version 2.0) comprises de-identified patient data, including characteristics like age, gender, ethnicity, and marital status, extracted from individuals who were admitted to Beth Israel Deaconess Medical Center (BIDMC). Furthermore, MIMIC-IV-ED (version 2.2) is an extensive database of emergency department (ED) admissions at the BIDMC between 2011 and 2019, which contains detailed clinical information, including diagnosis, medication, triage, and vital signs.

Each of the databases comprises distinct tables containing varying details related to a patient's hospitalisation. An individual patient is assigned a unique identifier, referred to as the subject id. However, since a single patient might have multiple hospitalisations, or a single stay may generate several records, linking these databases using subject id proved unfeasible. Moreover, as the aim is to generate an accurate report, it is imperative that non-imaging data be collected within the same time frame as the chest x-ray. Consequently, we resorted to record linkage between MIMIC-CXR and MIMIC-IV-ED databases and extracted data only if the patient was in the ED while the report was being generated and did not leave during that period. After performing data cleaning procedures; excluding missing data, filtering the images to only include anteroposterior and posteroanterior projections, keeping only one study if there is more than one records, and removing replications, the resulting dataset contains 65813 entries and 11 features including acuity level, oxygen saturation, heart rate, respiratory rate, systolic blood pressure, diastolic blood pressure, temperature, patient's chief complaint, ICD title, gender and ethnicity.

One challenge with the dataset for this task is its biases, particularly, its skewed distribution towards normal cases and the presence of numerous identical reports for different patients. To minimise these issues, we selected a subset of 65813 entries, by identifying and cataloging unique medical reports, ensuring that each distinct group was represented in the curated dataset to a similar extent. This subset consisted of 3000 total samples, which we further divided into a training set comprising 2100 datapoints and a validation set comprising 900 datapoints. Subsequently, we evaluated the performance of our models on a holdout test set comprising 1173 unseen examples. As there is currently no comparable comprehensive dataset encompassing similar non-imaging data, we exclusively employed this specific dataset to train and assess the proposed approach.

3.2 Feature extraction and pre-processing

This section describes the pre-processing and encoding of different data modalities used in this study. The main objective is to bridge the gap between the data used in the study and the data typically encountered in medical practice while minimising potential biases that may arise.

3.2.1 Image data

Each image went through resizing to 299 pixels \times 299 pixels, followed by min-max normalisation to scale the intensity values to a range of 0 to 1. The process of obtaining the representation of each image can be described as a two-step procedure. In the first step, the EfficientNet model is utilised as the base model to extract the visual features of the image. In the second step, this feature vector is employed as the input for a transformer-based encoder which extracts higher-level features and fuses this information with clinical and non-clinical (demographic) data. A detailed explanation of the fusion process can be found in [Section 3.3](#).

3.2.2 Clinical: non-imaging data

This study exclusively employed clinical data that clinicians considered during patient evaluations, wherein, a chest x-ray examination was conducted if any disease/abnormalities were suspected. These data included heart rate, respiratory rate, oxygen saturation, temperature, level of acuity (severity), primary symptoms or complaints, as well as known or suspected diseases.

The acuity level of a patient is determined based on the triage assessment, and an integer value, *between 1 and 5*, is assigned to each case where 1 indicates the least severe and 5 is the most severe. The higher acuity levels are typically associated with the presence of abnormalities in the patient's case, therefore, utilising the acuity level may assist the network in determining normal and abnormal cases while generating the report.

The integer-based variables including oxygen saturation, heart rate, respiratory rate, systolic blood pressure (sbp), and diastolic blood pressure (dbp) were initially treated to remove outliers. Subsequently, the values have been normalised within the range of 0 to 1, based on their respective minimum and maximum values. As for temperature data, a conversion to Fahrenheit scale

was performed, and similar to the integer-based variables, the values were normalised between 0 and 1.

The text-based variables, namely the chief complaint and ICD title variables, are initially processed by converting characters to lowercase and removing unnecessary punctuation, such as commas, periods, and newline characters, utilising regular expressions. Consecutive periods are condensed into single spaces, and double periods are substituted with single spaces, contributing to a more consistent text format. The resultant text undergoes further standardisation by substituting shorthand phrases or abbreviations with their corresponding full-text counterparts. For example, "cp" is replaced with "chest pain", "sob" or "shortness of breath" is replaced with "dyspnea" and so on. Standardisation also includes converting phrases like "chest pain, dyspnea" into "chest pain and dyspnea" as well as fixing typos and pluralisation issues such as changing "fevers" to "fever."

3.2.3 Non-clinical data

In addition to clinical data, patient records often include non-clinical metadata that can provide valuable insights. This study concentrates on two commonly collected non-clinical variables: gender, and ethnicity. These variables have been demonstrated to have an impact on health outcomes (30) and are therefore of particular interest in this paper.

As the gender data is already in binary format, the only necessary pre-processing step was to convert the data to a numerical representation by replacing "Male" with 0 and "Female" with 1. The ethnicity data was initially categorised into 5 broad groups consisting of the most frequently occurring values and this initial categorisation slightly improved model performance. The data was then categorised in a more granular fashion into 9 groups: White, African American, Hispanic/Latino, Black, Asian, White/European, Russian, Other, and Unknown. We hypothesised that employing these more detailed ethnicity categories would enable more accurate report generation. Subsequently, the categorical ethnicity data was mapped to integer values and reshaped into a 2D array to allow for input into the encoder.

3.3 Multi-modal data fusion details

This section provides the details of the multi-modal data fusion strategy utilised in our proposed model. As shown in [Figure 2](#), our approach employs a cross-attention mechanism to fuse the textual, visual, and scalar modalities. The key variables used in the fusion are defined in [Table 1](#). Specifically, the scalar patient data, comprising attributes like heart rate, oxygen saturation, respiratory rate, blood pressure, temperature, acuity level and gender, is concatenated to form a continuous representation. This continuous data is then passed through a dense layer to produce a scalar output.

The scalar patient data, comprising attributes like heart rate, oxygen saturation, respiratory rate, blood pressure, temperature, acuity level and gender, is concatenated to form a continuous representation. This continuous data is then passed through a dense layer to produce a scalar output.

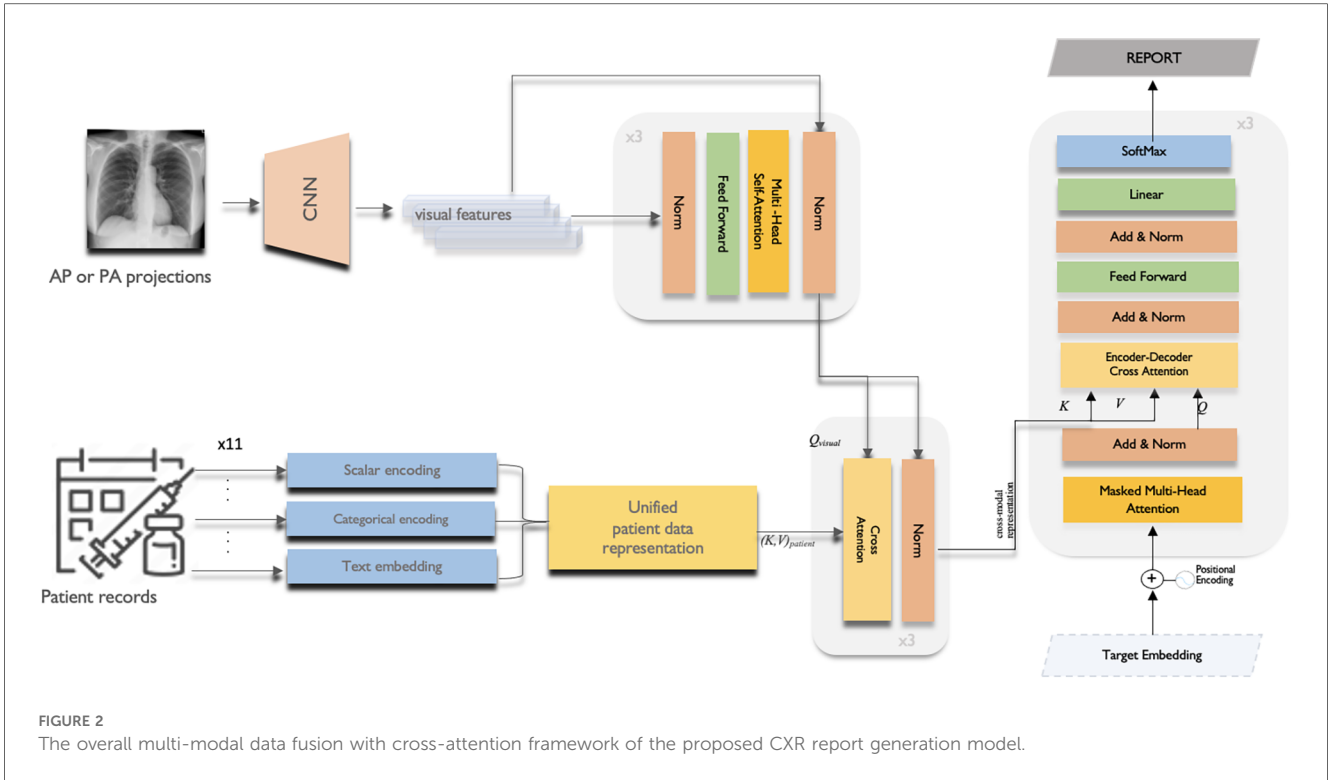


TABLE 1 List of variables.

Variable	Description
X_{text_data}	2D input tensor of text data indices
V_{text_data}	Vocabulary size of text data
E	Embedding dimension
W_{emb}	Embedding weight matrix
f_{embed}	Embedding function
X_{chief_embed}	Embedded chief complaint data
X_{icd_embed}	Embedded ICD title data
X_{scalar}	Processed scalar patient data
X_{eth}	One-hot encoded ethnicity data
$X_{patient}$	Unified patient representation
Q	Query matrix for attention
K	Key matrix for attention
V	Value matrix for attention

Each ethnicity group variable is transformed using the one-hot encoding (Equation 1), resulting in a matrix where each individual’s ethnicity is represented as a binary vector.

$$X_{eth} = [\delta(eth, 1), \delta(eth, 2), \dots, \delta(eth, 9)]. \quad (1)$$

where the function $\delta(i, j)$ is defined as

$$\delta(i, j) = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j. \end{cases}$$

The chief complaint and ICD title data consist of text sequences with varying lengths and vocabulary sizes. Therefore, these data

are separately embedded using the following embedding technique before being further processed through dense layers.

Let,

$$x_{text_data} \in \mathbb{Z}^{N \times M} \text{ — 2D input tensor of indices}$$

$$V_{text_data} \text{ — vocabulary size}$$

$$E \text{ — embedding dimension}$$

$$W_{emb} \in \mathbb{R}^{V_{text_data} \times E} \text{ — embedding weight matrix}$$

Then,

$$f_{embed}(x_{text_data}) \in \mathbb{R}^{N \times M \times E} \quad (2)$$

$$f_{embed}(x_{text_data})_{i,j,k} = W_{emb}[x_{text_data}_{i,j}, k]$$

where

$$f_{embed}(x_{text_data})_{i,j,k} \text{ — embedding vector}$$

$$\text{for token at position } (i, j)$$

$$x_{text_data}_{i,j} \text{ — integer index of token at position } (i, j)$$

$$W_{emb}[x_{text_data}_{i,j}, k] \text{ — } k\text{-th value from row of}$$

$$W_{emb} \text{ for index } x_{text_data}_{i,j}$$

$$X_{chief_embed} = f_{embed}(x_{chief_data})$$

$$X_{icd_embed} = f_{embed}(x_{icd_data})$$

where: x_{chief_data} and x_{icd_data} are the respective input indices tensors and f_{embed} is the embedding function defined in Equation 2.

After feature extraction and transformation of patient data inputs, the representations are concatenated into a unified patient representation vector (Equation 3). The processed scalar data output $X_{\text{scalar}} \in \mathbb{R}^{N \times M_{\text{scalar}}}$, one-hot encoded ethnicity output $X_{\text{eth}} \in \mathbb{R}^{N \times M_{\text{eth}}}$, and embedded chief complaint and ICD title outputs $X_{\text{chief_embed}}, X_{\text{icd_embed}} \in \mathbb{R}^{N \times M \times E}$ are concatenated for each patient giving:

$$X_{\text{patient}} = \text{Concatenate}(X_{\text{scalar}}, X_{\text{eth}}, X_{\text{chief_embed}}, X_{\text{icd_embed}}; \text{axis} = 1) \quad (3)$$

Where the Concatenate() operation joins the input tensors along the specified dimension, in this case, axis=1, yielding:

$$X_{\text{patient}} = \begin{bmatrix} X_{\text{scalar}} & X_{\text{eth}} & X_{\text{chief_embed}} & X_{\text{icd_embed}} \end{bmatrix} \in \mathbb{R}^{N \times (M_{\text{scalar}} + M_{\text{eth}} + 2M \times E)} \quad (4)$$

The resulting X_{patient} (Equation 4) contains a unified representation of each patient's data for further use, combining structured scalar variables, categorical encodings, and semantically rich embedded features into a single vector. This concatenation enables the joint modeling of heterogeneous data types into an integrated patient representation.

Then, an EfficientNetB0 CNN backbone (31) pre-trained on ImageNet extracts features from 299x299x3 RGB input images. The CNN outputs $N \times D$ image embeddings, where N is batch size and D is the feature dimension. This 1280-length visual feature vector is transformed via layer normalisation and a dense layer to refine the image representation. Before starting to fusion operation, multi-headed self-attention (Equation 5 and 6) is then applied to enable the model to jointly focus on different positions in the image via parallel heads. The self-attention outputs are then added to the original embedded image via residual connection, and normalised by a layernorm layer.

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (5)$$

where $\sqrt{d_k}$ is the dimension of the key vector k and query vector q

$$\text{MultiHead}(Q, K, V) = \text{Concat}(\text{head}_1, \dots, \text{head}_h)W^O \quad (6)$$

where

$$\text{head}_i = \text{Attention}(QW_i^Q, KW_i^K, VW_i^V)$$

The final output image embedding is further contextualised with information from the entire set of patient data via a cross-attention mechanism. In the cross-attention module, the convention is to take the image features as the query (Q) and the unified patient representation X_{patient} as the key (K) and value (V). This allows each part of the encoded image embedding to

attend to relevant semantics from the full patient data:

$$Q = \text{Image Features} \in \mathbb{R}^{N \times D}$$

$$K = V = X_{\text{patient}} \in \mathbb{R}^{N \times D}$$

Where N is the batch size and D is the common embedding dimension across modalities.

Multi-headed scaled dot-product attention is again applied between Q and K to obtain attention weights representing the relevance of each part of the patient data to each part of the image. The weighted value matrices are concatenated and projected to obtain the cross-attention outputs allowing the model to condition each part of the image embedding on relevant unified patient representation.

The cross-attention outputs are residually connected and normalised in a similar manner via element-wise addition with the output image embedding from the previous self-attention block and layernorm.

We adopt the canonical Transformer decoder architecture which starts by embedding the input sequence using both target (token) embeddings and positional embeddings. Target embeddings provides the meaning of words, while positional embeddings provide information about the order of tokens in the sequence. The initial layer employs self-attention, where each output token attends over previously generated tokens. This auto-regressive nature allows the model to condition on its own past predictions. Next, it performs attention over the encoded cross-modal representation obtained from the encoder. This enables each decoded output embedding to be conditioned on the relevant semantic concepts and modalities from the encoder, facilitating more effective fusion and reasoning across modalities.

4 Experiments and analysis

4.1 Experimental setup

The model undergoes training through a custom loop that involves the following key steps: data retrieval, image embedding, encoding of clinical and non-clinical data, calculation of loss and accuracy, computation of gradients, weight updating, and tracker adjustment.

The standardised sequence lengths are 43 tokens for reports, 2 tokens for chief complaints, and 6 tokens for ICD codes. The vocabularies contain over 6,000 unique tokens for radiology reports and over 3,000 tokens each for chief complaints and ICD codes. The vocabulary size and fixed sequence length were determined based on the complete dataset, not just the balanced subset of 3,000 samples used for training and validation. Both image features and text tokens are represented using 512-dimensional embeddings. The Transformer encoder and decoder layers include feed-forward networks with 512-dimensional units each, and Transformer layers utilize multi-headed attention with 3 attention heads. During training, a batch size of 64 is

employed, and training proceeds for 100 epochs with early stopping triggered by validation loss stagnation over 5 epochs.

The model's training employed the Adam optimiser with a learning rate of $3e-4$ and linear warmup for the first 500 steps. After the warm-up phase, the learning rate remains constant, stabilising training and facilitating effective model fine-tuning. Loss is calculated using the Sparse Categorical Cross-Entropy loss function defined in Equation 7, and accuracy is assessed by matching predicted tokens with true tokens. Let: y_{true_i} be the ground truth for the radiology reports. y_{pred_i} be the output from our report generation model. The equation for cross-entropy loss for each report without reduction is given by:

$$\text{loss}_i = - \sum (y_{true_i} \cdot \log((y_{pred_i}))) \quad (7)$$

where: i represents the index of the report. y_{true_i} is the ground truth for report i . y_{pred_i} is the generated report for index i . This loss calculation is performed for each report separately, without any reduction.

4.2 Evaluation metrics

To evaluate the linguistic quality of the generated radiology reports, we computed several automatic evaluation metrics comparing the generated text to the reference reports. First, BLEU-1 to BLEU-4 scores (32) were calculated to assess n-gram precision for unigrams up to 4-grams. Higher BLEU scores indicate greater local word-level similarity between the generated and reference texts. Second, the ROUGE-L score was used to measure the longest common subsequence, assessing the quality of the generated text in terms of recall and precision.

Additionally, we evaluated semantic similarity using the BERT Score (33). However, generic BERT representations may not fully capture domain-specific conceptual information needed for clinical text generation. To address this, we initialise BERTScore with the Bio-ClinicalBERT embeddings (34) trained on scientific text and clinical notes. By plugging these contextual embeddings into the BERTScore framework, we obtain a domain-adapted evaluation metric that emphasises clinical conceptual similarity. Specifically, the F1 component now computes semantic textual similarity using Bio-ClinicalBERT's clinical embeddings rather than generic BERT. These metrics providing a more nuanced assessment of meaning compared to strict n-gram matching. The BERT-based metrics can capture whether the generated reports convey clinically coherent descriptions despite differing word usage compared to the reference. Collectively, these automated evaluation metrics quantify linguistic similarity at both word-level, sentence-level, and semantic meaning levels.

4.3 Quantitative results

In this study, we leverage 11 additional clinical features along with chest x-rays to generate more accurate and patient-informed

radiology reports. The baseline model only employed chest x-ray images as input to generate corresponding reports, serving as our benchmark reference where the sole source of information was the visual data.

In order to analyse the contribution of each distinct data feature to model performance, we conducted an ablation study by incrementally presenting different features alongside the chest x-ray images.

For a fair comparison, all data features were encoded in the same way across all experiments, and model hyperparameters, as well as dataset splits, remained consistent.

We evaluated four main approaches:

1. The singular model incorporated a single additional feature to show individual performance. Oxygen saturation (O2Sat) is chosen for comparison as it demonstrated the highest performance among the singular models, as illustrated in Table 2.
2. The TextFusion model explored fusing textual features of reported primary symptoms and ICD diagnostic codes with the chest x-rays.
3. The ScalarFusion approach combined multiple predictive scalar features with the images, including O2Sat, diastolic blood pressure, temperature, patient acuity scores, and gender. Each of these scalars individually demonstrated performance improvements in singular models.
4. At the core of our study, the FullFusion model takes a holistic approach by fusing all available and relevant data points. This multi-modal fusion aims to effectively incorporate the diverse sources of information at hand, including chest x-ray images, structured clinical data, and unstructured text notes.

Table 3 presents a quantitative comparison based on the performance across multiple evaluation metrics. The metrics utilised for the assessment include BLEU-n (B_1 to B_4), ROUGE-L (R_L), BERT F1Score (BS_{F1}), and Bio-ClinicalBERT F1Score (Bio-CBS $_{F1}$). The SingularO2Sat method displayed notable improvements across multiple metrics compared to the baseline, while the TextFusion and ScalarFusion methods showcased marginal increases. The FullFusion method emerged as the top performer, showing substantial enhancements in various metrics, and highlighting the benefits of multi-modal fusion.

TABLE 2 Singular model results.

Singular Model	B_1	B_2	BS $_{F1}$	Bio-CBS $_{F1}$	R_L
ETHNICITY	0.328	0.212	0.174	0.782	0.321
HEARTRATE	0.333	0.213	0.170	0.786	0.295
ICD	0.328	0.207	0.186	0.785	0.301
RESPRATE	0.329	0.213	0.185	0.788	0.309
SBP	0.336	0.219	0.197	0.786	0.319
DBP	0.338	0.220	0.188	0.784	0.322
O2SAT	0.343	0.222	0.199	0.789	0.321
ACUITY	0.342	0.222	0.200	0.789	0.309
TEMPERATURE	0.336	0.219	0.199	0.789	0.326
GENDER	0.341	0.224	0.197	0.787	0.331
CHIEFCOMPLAINT	0.326	0.212	0.185	0.785	0.302

The highest-performing results are highlighted in bold in the table.

TABLE 3 Quantitative comparison of fusion methods: performance evaluation across multiple metrics. B_n for BLEU-n, R_L for ROUGE-L, BS_{F1} for BERT Score F1Score and CBS_{F1} for Bio-ClinicalBERT Score F1Score.

Method	B_1	B_2	B_3	B_4	R_L	BS _{F1}	Bio-CBS _{F1}
Baseline	0.326	0.205	0.138	0.084	0.301	0.192	0.787
Singular02Sat	0.343	0.222	0.151	0.096	0.321	0.199	0.789
TextFusion	0.326	0.209	0.141	0.862	0.307	0.181	0.784
ScalarFusion	0.343	0.219	0.145	0.090	0.320	0.198	0.786
FullFusion	0.351	0.231	0.162	0.107	0.331	0.218	0.794

The highest-performing results are highlighted in bold in the table.

Among the test samples, approximately 33% of them have BLEU-1 scores between 0.1 and 0.3, around 54% have scores between 0.3 and 0.5, about 11% have scores between 0.5 and 0.7, and a mere 0.26% have scores between 0.7 and 1, indicating high similarity. Our BLEU-1 results exhibit strong concordance with existing report generation literature, which has established scoring norms averaging 0.3 to 0.4 for this metric. We conducted a comparative analysis of our model against relevant state-of-the-art models (23, 24, 7, 22), referencing the results documented in their published literature. When considering ROUGE-L score, which reflects the model's ability to capture document-level linguistic coherence, our approach excelled in this aspect, achieving a ROUGE-L score of 0.331, standing out as the highest score across all models. This suggests that our model excels in capturing the long-range linguistic context of medical reports.

4.4 Qualitative results

For better interpretation of the results, we illustrated the samples in Figure 3 that showcase of diversity such as accurate prediction, different expressions, missing and false arguments, and completely false prediction. We compared the ground truth with our FullFusion model, and the correctly predicted diagnoses highlighted in bold for emphasis.

The results show promise in producing reports that capture many of the key findings described in the ground truth reports. To begin with, in all cases, the order of findings aligns with the reports written by the radiologists and the generated reports are structurally correct. The results also reveal a generally positive alignment in terms of language and grammar, however, some of the generated reports exhibit repeated words or phrases, which can affect the overall coherence. Additionally, the usage of "and" at the beginning of sentences and concluding paragraphs with "the" or "is" reveals grammatical inconsistencies. Furthermore, the FullFusion model accurately identifies normal cardiac, mediastinal, and hilar contours when present in the ground truth. It also reliably notes the presence or absence of abnormalities like pulmonary edema, pleural effusion, focal consolidation, pneumonia, and pneumothorax which are crucial in radiology analysis. In some cases, the generated reports exhibit a reduced level of detail compared to the ground truth, omitting certain specific observations.

In the first sample, the model missed the right middle lobe atelectasis that was noted in the ground truth. In the sample 2, the model hallucinated mediastinal clips not present in the

ground truth or image. The sample 3 shows that the model did not fully capture the enlarged cardiac silhouette and vessels described in the ground truth. In the sample 4, the model missed details about the interval removal of a central venous line and differences in positioning compared to a prior exam that provided important clinical context. In the sample 5, the model demonstrated enhanced detail compared to the ground truth by providing additional descriptive findings.

Sample 6 shows that the model failed to identify the surgical clips in the right upper quadrant indicating a prior cholecystectomy that were noted in the ground truth. The model also incorrectly identified findings suggestive of chronic obstructive pulmonary edema in the upper quadrant. Sample 7 had repetitive phrasing about no acute osseous abnormalities and failed to note the subsegmental atelectasis in the left lung base documented in the ground truth. Otherwise, the report accurately stated that the heart size was normal, the lungs were clear, and no effusions or pneumothorax were present. Lastly, in the sample 8, the model did not fully capture the moderate cardiac enlargement and aortic tortuosity described in the ground truth, instead stating mild cardiac enlargement. The predicted report also repeats "stable mediastinal silhouettes are stable" incorrectly. However, it accurately notes the lack of pulmonary effusion, pneumothorax or consolidation similar to the ground truth.

Overall, these qualitative results demonstrate good progress for the radiology report generation model, with accurate high-level identification of key findings, but also room for improvement in capturing more nuanced details and clinical context. While the baseline model was capable of providing results, it did not exhibit the same level of detail and accuracy as the enhanced model. Continued refinement of the model will be important to ensure accurate detection and description of abnormalities.

4.5 Radiologist evaluation results

We evaluated the model using 158 randomly selected samples from the unseen test set, covering diverse medical conditions reflecting the full distribution. A board-certified radiologist assessed three criteria: language fluency, content selection, and correctness of abnormal findings (AF) (Table 4). For language, the radiologist evaluated sentence structure, terminology, and overall clarity. For content, they compared the report's level of detail, key findings, and image coverage to the true findings.

Input Image				
	Ground Truth	<p>cardiac mediastinal and hilar contours are normal the pulmonary vasculature is normal minimal atelectasis is seen in the right middle lobe no focal consolidation pleural effusion or pneumothorax is present there are no acute osseous abnormalities</p>	<p>pa and lateral views of the chest were obtained demonstrating clear lungs without focal consolidation effusion or pneumothorax cardiomeastinal silhouette appears normal the imaged osseous structures are intact no free air below the right hemidiaphragm</p>	<p>enlargement of the cardiac silhouette and increase caliber pulmonary vessels is in keeping with the patients pregnant status allowing for low lung volumes which results in crowding of bronchovascular structures there is no evidence of pulmonary edema or focal consolidation to suggest pneumonia no definite pleural effusion</p>
Generated Report		<p>mediastinal and hilar contours are normal the pulmonary vasculature is normal streaky opacities in the lung bases likely reflect atelectasis no focal consolidation pleural effusion or pneumothorax is seen there are no acute osseous abnormalities the left upper abdomen is unremarkable there</p>	<p>and mediastinal clips are noted in the right upper lobe there is no focal consolidation effusion or pneumothorax the cardiomeastinal silhouette is normal imaged osseous structures are intact no free air below the right hemidiaphragm is seen in the right hemidiaphragm is</p>	<p>enlargement of the cardiac silhouette is mildly enlarged the mediastinal and hilar contours are unremarkable pulmonary vasculature is normal lungs are clear with no pleural effusion or pneumothorax is seen there are no acute osseous abnormalities are again noted in the right</p>
	Input Image			
Ground Truth		<p>the cardiomeastinal silhouette is normal there is no focal consolidation there is no pleural effusion or pneumothorax</p>	<p>the lungs are clear the cardiomeastinal silhouette is normal surgical clips in the right upper quadrant suggest prior cholecystectomy</p>	<p>cardiac mediastinal and hilar contours are unchanged with the heart size normal the pulmonary vasculature is not engorged re demonstrated within the left lung base is subsegmental atelectasis without focal consolidation the right lung is clear no pleural effusion or pneumothorax is present no acute osseous abnormality is detected</p>
	Generated Report	<p>there is no focal consolidation pleural effusion or pneumothorax the cardiomeastinal silhouette is normal there is no acute osseous abnormalities are seen there is no free air below the right hemidiaphragm is seen in the right</p>	<p>chest film demonstrates no focal consolidation pleural effusion or pneumothorax the cardiac and mediastinal silhouettes are unremarkable no acute osseous abnormalities are seen within the diaphragm suggesting chronic obstructive pulmonary edema is seen at the upper quadrant of uncertain chronicity pulmonary edema</p>	<p>tip of the heart size is normal cardiomeastinal contours are normal the pulmonary vasculature is normal lungs are clear without focal consolidation pleural effusion or pneumothorax is normal no acute osseous abnormalities identified no acute osseous abnormalities identified there is identified there</p>

FIGURE 3
GT and GR report from the proposed FullFusion CXR report generation model.

TABLE 4 Radiologist evaluation results on a 1–5 scale.

Language fluency	Content selection	Correctness of AF
4.24	4.12	3.89

They assessed accuracy of abnormal findings by comparing to the true conditions. The radiologist assigned 1–5 scores and noted preference between reports. This methodology enabled quantitative and qualitative assessment of language generation, content selection, and diagnostics. The radiologist also noted that while performing well overall, some shortcomings were observed. The model often missed surgical materials like catheters and clips and it fails to capture anomaly variations when the patient is inclined to the right or left. Sensitivity to bone lesions was lacking, overlooking non-urgent findings like scoliosis. However, it's worth noting that these are not extensively covered in the ground truth as well. For normal x-rays, it occasionally included non-definitive elements. While these additions may be accurate, there is a slight possibility that they may not be. This evaluation methodology provided valuable insights into model strengths and areas needing improvement.

5 Discussion and conclusion

In this study, we adopted a comprehensive approach to enhance the precision and clinical relevance of radiology reports generated in conjunction with chest x-ray images. We achieved this by incorporating cutting-edge network components and drawing inspiration from state-of-the-art methods, such as the transformer architecture and multi-modal data fusion techniques. In the section pertaining to data selection, we ensured the inclusion of data only within the time frame of radiology report generation. This was done to closely replicate the clinical pathway. Additionally, we aimed for a balanced representation of each type of report in our sample selection to counteract potential biases skewed towards normal cases. Although this approach resulted in a smaller dataset compared to existing literature, it was a crucial step for preventing biased results and validating the results in real clinical settings.

Recent literature highlights the potential of multi-modal learning techniques in advancing the quality of automated radiology report generation. However, a majority of medical report generation models primarily focus on target reports within specific information or incorporate image findings as supplementary inputs. Given that Chest x-rays present three-dimensional objects in a two-dimensional form, some valuable information is lost, leading to semantic gaps in the data provided to the network. Furthermore, radiologists possess more data beyond images during report creation. To address this limitation, we bridged the semantic gap between vision and language models by capturing uncodified information essential to the diagnosis process. We achieved this by introducing an ensemble of 11 supplementary features in conjunction with the chest x-ray data. These features were thoughtfully selected to enhance both accuracy and clinical insight in the generated reports.

The results indicate that incorporating non-imaging clinical and non-clinical data positively impacts the quality of the generated reports. Our ablation study further demonstrates that providing all data simultaneously yields higher accuracy compared to using individual data components separately. This finding suggests that introducing data with no significant standalone impact on the network alongside others leads to improved results. In light of these outcomes, we recommend that the broader research community include additional metadata for enhanced model performance.

However, there are some limitations to our study. While the multi-modal deep neural network framework holds potential strength, its complexity and resource-intensive nature may pose challenges. This might hinder its real-time application in medical settings, especially those with limited resources and computational power. Furthermore, our data solely originates from databases within a single institution, lacking a comparable comprehensive dataset that combines imaging and non-imaging data (both clinical and non-clinical) with linked radiology reports. Enhancing data diversity from various sources could enhance the overall robustness of the study.

Data availability statement

Publicly available datasets were analyzed in this study. This data can be found here: <https://physionet.org/content/mimic-cxr>.

Ethics statement

Ethical approval was not required for the study involving humans in accordance with the local legislation and institutional requirements. Written informed consent to participate in this study was not required from the participants or the participants' legal guardians/next of kin in accordance with the national legislation and the institutional requirements.

Author contributions

NA: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Resources, Software, Validation, Visualization, Writing – original draft, Writing – review & editing; SS: Writing – review & editing, Methodology, Project administration, Supervision; SB: Writing – review & editing, Validation; NR: Supervision, Writing – review & editing, Conceptualization, Methodology, Project administration; AFF: Writing – review & editing, Conceptualization.

Funding

The author(s) declare financial support was received for the research, authorship, and/or publication of this article.

This work was supported by Milli Eğitim Bakanlığı and University of Leeds.

Acknowledgments

This study is fully sponsored by the Turkish Ministry of National Education.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial

relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

- Yang S, Niu J, Wu J, Liu X. Automatic medical image report generation with multi-view, multi-modal attention mechanism. In: *International Conference on Algorithms, Architectures for Parallel Processing*. Springer (2020). p. 687–99.
- Singh S, Karimi S, Ho-Shon K, Hamey L. Show, tell and summarise: learning to generate and summarise radiology findings from medical images. *Neural Comput Appl*. (2021) 33:7441–65. doi: 10.1007/s00521-021-05943-6
- Jing B, Xie P, Xing E. On the automatic generation of medical imaging reports. *arXiv [Preprint] arXiv:1711.08195* (2017).
- Xue Y, Xu T, Rodney Long L, Xue Z, Antani S, Thoma GR, et al. Multimodal recurrent model with attention for automated radiology report generation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer (2018). p. 457–66.
- Yuan J, Liao H, Luo R, Luo J. Automatic radiology report generation based on multi-view image fusion and medical concept enrichment. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer (2019). p. 721–9.
- Chen Z, Song Y, Chang TH, Wan X. Generating radiology reports via memory-driven transformer. *arXiv [Preprint] arXiv:2010.16056* (2020).
- Nooralhazadeh F, Gonzalez NP, Frauenfelder T, Fujimoto K, Krauthammer M. Progressive transformer-based generation of radiology reports. *arXiv [Preprint] arXiv:2102.09777* (2021).
- You Q, Jin H, Wang Z, Fang C, Luo J. Image captioning with semantic attention. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2016). p. 4651–9.
- Anderson P, He X, Buehler C, Teney D, Johnson M, Gould S, et al. Bottom-up and top-down attention for image captioning and visual question answering. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2018). p. 6077–86.
- Ghandi T, Pourreza H, Mahyar H. Deep learning approaches on image captioning: A review. *ACM Comput Surv*. (2023) 56:1–39. doi: 10.1145/3617592
- Suresh KR, Jarapala A, Sudeep P. sImage captioning encoder-decoder models using CNN-RNN architectures: A comparative study. *Circ Syst Signal Process*. (2022) 41:5719–42. doi: 10.1007/s00034-022-02050-2
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. *Adv Neural Inf Process Syst*. (2017) 30:6000–10. doi: 10.48550/arXiv.1706.03762
- Zhu X, Li L, Liu J, Peng H, Niu X. Captioning transformer with stacked attention modules. *Appl Sci*. (2018) 8:739. doi: 10.3390/app8050739
- Xie S, Girschick R, Dollár P, Tu Z, He K. Aggregated residual transformations for deep neural networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2017). p. 1492–500.
- Zhang W, Nie W, Li X, Yu Y. Image caption generation with adaptive transformer. In: *2019 34th Youth Academic Annual Conference of Chinese Association of Automation (YAC)*. IEEE (2019), p. 521–6.
- Li G, Zhu L, Liu P, Yang Y. Entangled transformer for image captioning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2019). p. 8928–37.
- Cornia M, Stefanini M, Baraldi L, Cucchiara R. Meshed-memory transformer for image captioning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020). p. 10578–87.
- Liu W, Chen S, Guo L, Zhu X, Liu J. CPTR: Full transformer network for image captioning. *arXiv [Preprint] arXiv:2101.10804* (2021).
- Irvin J, Rajpurkar P, Ko M, Yu Y, Ciurea-Ilcus S, Chute C, et al. CheXpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In: *Proceedings of the AAAI Conference on Artificial Intelligence* (2019), Vol. 33, p. 590–7.
- Xiong Y, Du B, Yan P. Reinforced transformer for medical image captioning. In: *International Workshop on Machine Learning in Medical Imaging*. Springer (2019). p. 673–80.
- Lewis M, Liu Y, Goyal N, Ghazvininejad M, Mohamed A, Levy O, et al. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv [Preprint] arXiv:1910.13461* (2019).
- Wang L, Ning M, Lu D, Wei D, Zheng Y, Chen J. An inclusive task-aware framework for radiology report generation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer (2022). p. 568–77.
- Yang S, Wu X, Ge S, Zhou SK, Xiao L. Knowledge matters: Chest radiology report generation with general and specific knowledge. *Med Image Anal*. (2022) 80:102510. doi: 10.1016/j.media.2022.102510
- Wu X, Li J, Wang J, Qian Q. Multimodal contrastive learning for radiology report generation. *J Ambient Intell Humaniz Comput*. (2023) 14:11185–94. doi: 10.1007/s12652-022-04398-4
- Stahlschmidt SR, Ulfenborg B, Synnergren J. Multimodal deep learning for biomedical data fusion: a review. *Brief Bioinform*. (2022) 23:bbab569. doi: 10.1093/bib/bbab569
- Wang X, Peng Y, Lu L, Lu Z, Summers RM. TieNet: Text-image embedding network for common thorax disease classification and reporting in chest x-rays. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2018). p. 9049–58.
- Chauhan G, Liao R, Wells W, Andreas J, Wang X, Berkowitz S, et al. Joint modeling of chest radiographs and radiology reports for pulmonary edema assessment. In: *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part II 23*. Springer (2020). p. 529–39.
- Hayat N, Geras KJ, Shamout FE. MedFuse: Multi-modal fusion with clinical time-series data and chest x-ray images. *arXiv [Preprint] arXiv:2207.07027* (2022).
- Lin CY. Rouge: A package for automatic evaluation of summaries. In: *Text Summarization Branches Out* (2004). p. 74–81.
- Aksoy N, Ravikumar N, Frangi AF. Radiology report generation using transformers conditioned with non-imaging data. In: *Medical Imaging 2023: Imaging Informatics for Healthcare, Research, and Applications*. SPIE (2023), Vol. 12469. p. 146–54.
- Tan M, Le Q. EfficientNet: Rethinking model scaling for convolutional neural networks. In: *International Conference on Machine Learning*. PMLR (2019). p. 6105–14.
- Papineni K, Roukos S, Ward T, Zhu WJ. BLEU: A method for automatic evaluation of machine translation. In: *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics* (2002). p. 311–8.
- Zhang T, Kishore V, Wu F, Weinberger KQ, Artzi Y. BERTScore: Evaluating text generation with bert. *arXiv [Preprint] arXiv:1904.09675* (2019).
- Alsentzer E, Murphy JR, Boag W, Weng WH, Jin D, Naumann T, et al. Publicly available clinical BERT embeddings. *arXiv [Preprint] arXiv:1904.03323* (2019).