Prior-RadGraphFormer: A Prior-Knowledge-Enhanced Transformer for Generating Radiology Graphs from X-Rays

Yiheng Xiong[1*], Jingsong Liu[1*], Kamilia Zaripova[1*], Sahand Sharifzadeh[2], Matthias Keicher[1†], and Nassir Navab[1†]

[1] Computer Aided Medical Procedures, Technische Universität München, Germany
[2] Ludwig Maximilians Universität München, Germany

Abstract. The extraction of structured clinical information from free-text radiology reports in the form of radiology graphs has been demonstrated to be a valuable approach for evaluating the clinical correctness of report-generation methods. However, the direct generation of radiology graphs from chest X-ray (CXR) images has not been attempted. To address this gap, we propose a novel approach called Prior-RadGraphFormer that utilizes a transformer model with prior knowledge in the form of a probabilistic knowledge graph (PKG) to generate radiology graphs directly from CXR images. The PKG models the statistical relationship between radiology entities, including anatomical structures and medical observations. This additional contextual information enhances the accuracy of entity and relation extraction. The generated radiology graphs can be applied to various downstream tasks, such as free-text or structured reports generation and multi-label classification of pathologies. Our approach represents a promising method for generating radiology graphs directly from CXR images, and has significant potential for improving medical image analysis and clinical decision-making.

Keywords: Radiology Graph Generation · Transformer · Prior Knowledge.

1 Introduction

In recent years, deep learning (DL) methods have significantly improved the computer-assisted diagnosis of chest X-ray (CXR) images. DL methods have been used in multiple ways. For instance, Ma and Lv [19] used a Swin transformer with fully-connected layers to classify CXR images as either normal or indicative of pneumonia. Cicero et al. [3] modeled the diagnosis process as a multi-label classification problem and used GoogleNet to classify CXR images into six categories. In comparison to these classification methods, generated free-text radiology reports can provide more comprehensive information about the impression and

---

[*] These authors contribute equally to this work and share first authorship.

[†] These authors share last authorship.

findings [9]. Shin et al.[23] pioneered CNN-RNN for automatic text generation. Wang et al.[25] introduced TieNet, generating reports and detecting thorax diseases. Hou et al.[6] proposed RATCHET, using transformer and attention to generate reports from CXR images. Kaur et al.[12] used contextual word representations for useful radiological reports. Cao et al.[1] proposed Multi-modal Memory Transformer Network for consistent medical reports. Existing approaches for automating medical reporting rely on generating free text, posing challenges for clinical evaluation [20,13]. To tackle this, ImaGenome [27] and RadGraph Benchmark[9] were introduced to extract structured clinical information from free-text radiology reports, represented as a radiology graph. Each node in a radiology graph corresponds a unique entity, such as an object with bounding boxes annotations or an attribute in ImaGenome or, an anatomical structure or an observation and its presence and uncertainty in RadGraph Benchmark. Although these graph representations of reports have been used to evaluate the clinical correctness of reports [28], generating radiology graphs directly from CXR images has not been explored. In contrast, diverse interactions between object pairs in natural images in the form of scene graph generation have been extensively investigated [17]. Li et al. [14], and Lu et al. [18] employed two-stage methods to propose dense relationships between predicted connected object pairs. In recent work, Shit et al. [24] introduced Relationformer, a unified one-stage framework based on DETR [2] that facilitates the end-to-end generation of graphs from images. It is a state-of-the-art method for detecting and generating graphs from natural images. However, it requires bounding boxes for the detected objects (nodes of the graph) and is not directly applicable to radiology graphs since some nodes (entities) in radiology graphs do not have exact locations, such as "left" or "clear", and most datasets do not provide bounding boxes annotations.

Therefore, we propose a detection-free method, Prior-RadGraphFormer, to generate radiology graphs directly from CXR images without requiring bounding boxes for each entity. The method incorporates prior knowledge in the form of probabilistic knowledge graphs (PKG) [22] that model the statistical relationship between anatomies and pathological observations. Experimental results show that Prior-RadGraphFormer achieves competitive results in the radiology-graph-generation task. Moreover, the generated graphs can be used for multiple downstream tasks such as generating free-text reports based on predefined rules, cheXpert labels [8] classification, and populating templates for structured reporting.

In summary, our contributions are the following: 1) proposing a novel detection-free method that generates radiology graphs directly from CXR images; 2) enhancing this method by incorporating prior knowledge, leading to improved performance; 3) extensively evaluating our method using RadGraph metrics and the two downstream tasks of report generation and multi-label classification of pathologies.
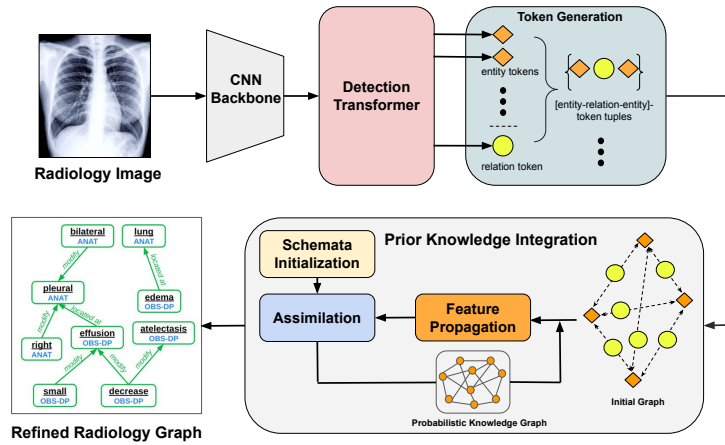
Fig. 1: **Prior-RadGraphFormer architecture.** A CNN backbone extracts radiology image features, which are fed into the detection transformer (DETR) [2] to generate multiple entity tokens and one relation token. Then an initial graph is constructed using entity and relation tokens. Each node is represented by a valid entity token and each edge is represented by a corresponding [entity-relation-entity]-token tuple. The initial graph is incorporated with prior knowledge and refined to output the final radiology graph. Our code is open sourced at https://github.com/xiongyiheng/Prior-RadGraphFormer

2   Method

We define radiology graph generation as the process of transforming CXR images into a graph structure introduced by RadGraph Benchmark [9] that represents the content of a radiology report describing the image. Each node in the graph corresponds to a unique entity, such as an anatomical structure or an observation, and its presence and uncertainty. The edges between nodes indicate relationships between these entities. A visual depiction of the proposed method is shown in Fig. 1. Prior-RadGraphFormer consists of two main components: RadGraphFormer and Prior Knowledge Integration.

2.1   RadGraphFormer

Shit et al. [24] proposed Relationformer, an image-to-graph framework that leverages direct set-based object prediction and incorporates the interaction among the objects to learn an object-relation representation jointly. Given an input image, Relationformer initially outputs a set of discrete object tokens ([obj]-tokens) and a relation token ([rln]-token) where the token embeddings come from applying cross attention between a set of pre-defined and randomly initialized token vectors to input image patches. Relationformer then predicts the corresponding class label and the location of the bounding boxes for each object. In addition,

Relationformer predicts a relation label for each pair of detected objects by concatenating each pair of [obj]-tokens with the [rln]-token and applying a relation prediction head on top.

In our model, we adapt Relationformer as the radiology graph generation backbone. We modify the entity prediction module by replacing the bounding box prediction head of it with the uncertainty prediction head due to the fact that radiology graphs do not contain bounding boxes but uncertainty information for each entity. Specifically, the entity prediction module comprises two distinct components. The first component is responsible for entity classification. The second component is responsible for predicting the uncertainty associated with each entity. Both are presented by one linear layer. To reflect these changes in terminology, we adjust the naming convention of the model's tokens, from [obj]-token to [ent]-token. Moreover, instead of using ResNet50 as the CNN backbone, we utilize DenseNet121 [7] based on its widespread adoption and effectiveness for processing CXR images [6,13].

2.2   Prior Knowledge Integration

Vanilla RadGraphFormer uses the concatenation of a pair of [ent]-tokens and a shared [rln]-token ($\{ent_i, r, ent_j\}_{i \neq j}$), followed by an MLP, for relation prediction. Although the transformer-based model inherently considers the context of the tokens, we argue that these representations may not adequately capture the complexity of the relationships. Consequently, this limited representation may lead to an incomplete or inaccurate understanding of the context, resulting in sub-optimal relation and/or entity prediction.

Following [22], in order to address this issue, we propagate higher-level prior knowledge in RadGraphFormer, as shown in the lower part of Fig. 1. To this end, we first construct an initial embedding graph G from all the valid ent_is as nodes. Each edge in the graph is represented by $edg_{ij} = MLP_{proj}(\{ent_i, r, ent_j\}_{i \neq j})$, where MLP_{proj} helps project the concatenated edge features to the same dimension of node features. This results in a fully-connected bi-directional graph. Then a stack of graph transformers GTs [29] are utilized to propagate features of both nodes and edges in G. In order to allow the storage and propagation of relational knowledge within our framework, we create a randomly initialized representation for each class as a trainable parameter s_c called schemata [22] that interacts with the outputs of the graph transformer. We then apply multiple assimilation steps [22]. The assimilation step is a process where the outputs of GT attend to s_cs such that the attention coefficients predict the classification outputs for each node/edge from the GT and the attention values are propagated to the output of GT. It is important to note that the attention coefficients are supervised by the ground truth labels for each entity/relation during the training. Details can be found in the appendix.

2.3   Training and Inference.

During training, we apply two supervisions for entity classes. One is for [ent]-tokens generated from DETR and the other is for the attention coefficients of nodes during each assimilation step. The latter one can be regarded as an additional entity class supervision. Entity uncertainty is supervised via [ent]-tokens and relation is supervised via the attention coefficients of edges during each assimilation step. During inference, we evaluate entity metrics solely based on [ent]-tokens. As for relation inference, we use the attention coefficients of edges from the last assimilation step as the classification output.

3   Experiments

3.1   Datasets

Our dataset comprises MIMIC-CXR-JPG v2.0.0 [4,11,10], which contains both imaging studies and free-text reports, and RadGraph Benchmark [9]. Our training set includes around 220,000 ground truth graph annotations obtained by RadGraph Benchmark, paired with the corresponding frontal CXR images. The validation set consists of 500 ground truth graph annotations obtained by board-certified radiologists and the corresponding frontal CXR images. There are 229 classes of entities after mapping, three levels of uncertainty, and three classes of relation. Details can be found in the appendix.

3.2   Implementation Details

All networks including the backbone are trained from scratch with PyTorch 1.12.0 and CUDA 11.6 on a single NVIDIA A40 until convergence. A batch size of 32 is chosen, and AdamW [16] optimizer with a learning rate of 1e-4 is utilized. To address the issue of imbalanced labels in entity class prediction, we employ focal loss [15], while cross entropy loss [5] is used for entity uncertainty prediction. For relation prediction, we utilize stochastic relation loss [24], with a foreground-to-background edge ratio of 1:3. The weights between entity classification loss, uncertainty loss, and relation loss are 1, 1, and 3. Additionally, two assimilation steps are performed during the prior knowledge integration process per iteration. CXR images are preprocessed following [13].

3.3   Evaluation

We evaluate our model similarly to RadGraph Benchmark [9] considering micro F1-score. A predicted entity is regarded as a true positive if its predicted entity class and uncertainty level are correct. For relation evaluation, a predicted relation is considered a true positive if its head entity, tail entity, and relation type are predicted correctly. To show the use cases for radiology graph generation in this work, we perform the evaluation for the two downstream tasks: free-text report generation and cheXpert labels [8] classification. Specifically, we generate free-text clinical reports from generated radiology graphs based on pre-defined rules.
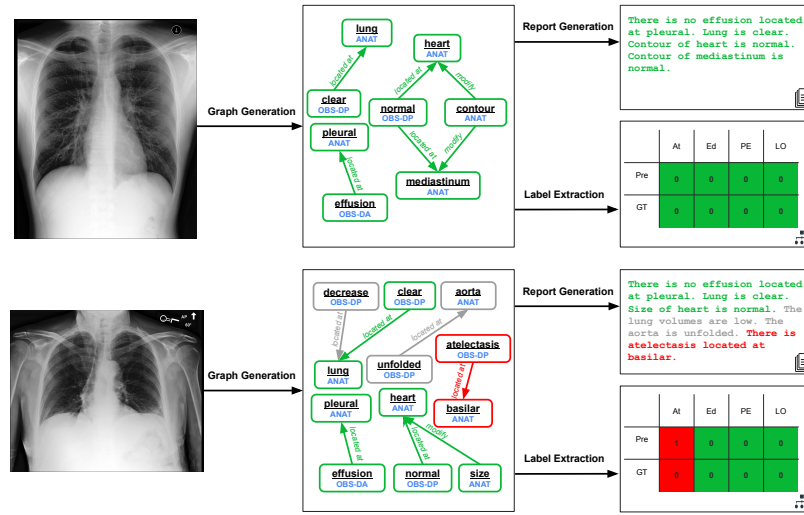
Fig. 2: **Qualitative results**. On the left and middle are the CXR images and the corresponding predicted radiology graphs by our model. On the right, we present the outputs of two downstream tasks derived via the predicted graph: free-text clinical report and pathologies classification result. True positives, false negatives, and false positives are represented by green, gray, and red, respectively.

4 Results and Discussion

4.1 Main Results

Given that no prior work has addressed our task, we create a baseline that can be evaluated via RadGraph metrics. The baseline is constructed by two separate pretrained models, where RATCHET [6] is used to predict free-text reports from CXR images and followed by RadGraph Benchmark [9] to produce radiology graphs from predicted reports. Note that radiology graphs generated from ground truth reports via pretrained RadGraph Benchmark are essentially ground truth labels as discussed in Section. 3.1 hence they are not directly comparable with our results. Additionally, we compare the performance of vanilla RadGraphFormer with Prior-RadGraphFormer. As is shown in Table. 1, both vanilla RadGraphFormer and Prior-RadGraphFormer demonstrate superior performance over the baseline in all RadGraph metrics. In particular, Prior-RadGraphFormer shows better performance than vanilla RadGraph-Former, indicating the significance of leveraging prior knowledge for radiology graphs generation.

4.2 Downstream Tasks

For downstream tasks, we compare our methods with RATCHET [6] using free-text NLP metrics and cheXpert label classification F1 score. However, please

Table 1: **Results comparisons on the validation set.** (P: precision, R: recall; B-1: BLEU-1, MET.: METEOR, R_L: ROUGE_L scores; At: Atelectasis, Ed: Edema, PE: Pleural Effusion, LO: Lung Opacity.) Best quality performance in bold.

Method	RadGraph Metrics						NLP Metrics				Classification F1			
	Entity			Relation			B-1	MET.	R_L	SPICE	At	Ed	PE	LO
	P	R	F1	P	R	F1								
RATCHET	-	-	-	-	-	-	**0.18**	**0.09**	0.19	**0.09**	0.38	0.27	0.47	0.17
RATCHET → RadGraph Benchmark	0.45	0.16	0.24	0.11	0.03	0.05	-	-	-	-	-	-	-	-
vanilla RadGraphFormer	0.58	0.47	0.52	0.35	0.19	0.25	0.14	0.08	**0.21**	0.08	**0.40**	0.26	**0.82**	**0.24**
Prior-RadGraphFormer	**0.63**	**0.55**	**0.59**	**0.44**	**0.21**	**0.28**	0.07	0.07	0.19	0.08	0.39	**0.29**	0.81	0.22

note that there may be variations in the numerical results of RATCHET due to differences in the validation set used in our experiments compared to the original paper. Table. 1 shows that our radiology graphs directly generated from CXR images can be properly transformed into free-text reports and pathologies classification results. They have reasonable performance in the corresponding evaluation and demonstrate the utility of a graph-based representation in the context of a clinical study.

From Table. 1 it is evident that vanilla RadGraphFormer exhibits better performance with respect to BLEU-1 score and F1 score of several pathologies as compared to Prior-RadGraphFormer. This disparity in performance can be explained by a more varied set of entities in radiology graphs generated by vanilla RadGraphFormer. Also, it is worth noting that the NLP metrics of RATCHET are generally better than our methods. However, higher NLP metrics may not indicate better clinical usefulness [20]. One example is shown in Table. 2. Qualitative results can be seen in Fig. 2.

Table 2: **Example of evaluating clinical reports with NLP metrics.** In this case, our report exhibits superior clinical accuracy compared to RATCHET, but NLP metrics associated with it are significantly lower.

Report section	NLP Metrics				Classification
	B-1	MET.	R_L	SPICE	PE
Ground Truth: As compared to the previous radiograph, small left and moderate layering right pleural effusions have increased in size.	-	-	-	-	Positive
RATCHET: As compared to the previous radiograph, the patient has been intubated.	0.36	0.19	0.47	0.24	Negative
Prior&Vanilla-RadGraphFormer: There is enlarged effusion located at bilateral of pleural.	0.07	0.06	0.13	0.12	Positive

Table 3: **Ablation studies on the validation set.** (AECS: additional entity class supervision)

Method	Entity			Relation		
	Precision	Recall	F1-score	Precision	Recall	F1-score
vanilla RadGraphFormer	0.58	0.47	0.52	0.35	0.19	0.25
Prior-RadGraphFormer w/o PKG	0.61	0.50	0.55	0.39	0.17	0.23
Prior-RadGraphFormer w/o AECS	0.63	0.49	0.55	0.43	0.17	0.24
Prior-RadGraphFormer	**0.63**	**0.55**	**0.59**	**0.44**	**0.21**	**0.28**

4.3   Ablation Studies

In our ablation studies, we aim to investigate two key aspects. Firstly, we aim to determine whether the integration of PKG truly improves performance or whether it is the graph transformers that enhance the model. Secondly, we aim to explore the impact of incorporating additional entity class supervision on entity and relation metrics.

To assess the impact of PKG integration, we simply use nodes features and edges features after graph transformers without assimilation to generate classification outputs. The results presented in Table. 3 demonstrate that Prior-RadGraphFormer without PKG is inferior to Prior-RadGraphFormer, thereby highlighting the importance of incorporating PKG as priors in Prior-RadGraph-Former.

As for the influence of incorporating additional entity class supervision in the assimilation step, we conduct a comparative experiment by removing it and analyzing the resulting performance. The experimental results displayed in Table. 3 show that the impact is positive.

4.4   Limitations and Outlook

In this work, we evaluated the performance without directly using the graph itself. In the future work we plan to investigate additional graph-specific metrics [21,26] to provide a more comprehensive evaluation. Furthermore, we acknowledge that the number of baselines is limited and aim to explore more comparative models in our ongoing research. Additionally, the inclusion of PKG increases the training time. In future work, we plan to optimize the integration of PKG to accelerate the training without compromising performance.

5   Conclusion

In this paper, we propose Prior-RadGraphFormer, a novel detection-free method that generates radiology graphs directly from CXR images. The model incorporates prior knowledge in the form of probabilistic knowledge graphs that capture the statistical relationship between anatomies and observations. We

demonstrate the effectiveness of Prior-RadGraphFormer in the CXR-image-to-radiology-graph task. Moreover, we show that generated radiology graphs are useful for downstream tasks such as free-text reports generation and multi-label classification of pathologies. Our findings offer a promising direction for automatically generating radiology graphs from CXR images. We pave the way to automate the classification of fine-grained clinical findings structured in a radiology graph that can be applied to generate reports while allowing for the assessment of clinical correctness.

References

1. Cao, Y., Cui, L., Zhang, L., Yu, F., Li, Z., Xu, Y.: Mmtn: Multi-modal memory transformer network for image-report consistent medical report generation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 277–285 (2023)
2. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16. pp. 213–229. Springer (2020)
3. Cicero, M.D., Bilbily, A., Colak, E., Dowdell, T., Gray, B., Perampaladas, K., Barfett, J.: Training and validating a deep convolutional neural network for computer-aided detection and classification of abnormalities on frontal chest radiographs. Investigative Radiology **52**, 281–287 (2017)
4. Goldberger, A.L., Amaral, L.A.N., Glass, L., Hausdorff, J.M., Ivanov, P.C., Mark, R.G., Mietus, J.E., Moody, G.B., Peng, C.K., Stanley, H.E.: Physiobank, physiotoolkit, and physionet: Components of a new research resource for complex physiologic signals. Circulation (2000)
5. Good, I.J.: Rational decisions. Journal of the Royal Statistical Society: Series B (Methodological) **14**(1), 107–114 (1952)
6. Hou, B., Kaissis, G., Summers, R.M., Kainz, B.: Ratchet: Medical transformer for chest x-ray diagnosis and reporting. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VII 24. pp. 293–303. Springer (2021)
7. Huang, G., Liu, Z., Pleiss, G., Van Der Maaten, L., Weinberger, K.: Convolutional networks with dense connectivity. IEEE Transactions on Pattern Analysis and Machine Intelligence (2019)
8. Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Ilcus, S., Chute, C., Marklund, H., Haghgoo, B., Ball, R., Shpanskaya, K., et al.: Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In: Proceedings of the AAAI conference on artificial intelligence. vol. 33, pp. 590–597 (2019)

9. Jain, S., Agrawal, A., Saporta, A., Truong, S.Q., Duong, D.N., Bui, T., Chambon, P., Zhang, Y., Lungren, M.P., Ng, A.Y., et al.: Radgraph: Extracting clinical entities and relations from radiology reports. arXiv preprint arXiv:2106.14463 (2021)
10. Johnson, A., Lungren, M., Peng, Y., Lu, Z., Mark, R., Berkowitz, S., Horng, S.: Mimic-cxr-jpg - chest radiographs with structured labels (version 2.0.0) (2019)
11. Johnson, A.E., Pollard, T.J., Greenbaum, N.R., Lungren, M.P., Deng, C.y., Peng, Y., Lu, Z., Mark, R.G., Berkowitz, S.J., Horng, S.: Mimic-cxr-jpg, a large publicly available database of labeled chest radiographs. arXiv preprint arXiv:1901.07042 (2019)
12. Kaur, N., Mittal, A.: Radiobert: A deep learning-based system for medical report generation from chest x-ray images using contextual embeddings. Journal of Biomedical Informatics **135**, 104220 (2022)
13. Keicher, M., Mullakaeva, K., Czempiel, T., Mach, K., Khakzar, A., Navab, N.: Few-shot structured radiology report generation using natural language prompts. arXiv preprint arXiv:2203.15723 (2022)
14. Li, R., Zhang, S., He, X.: Sgtr: End-to-end scene graph generation with transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 19486–19496 (June 2022)
15. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
16. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: International Conference on Learning Representations (2017)
17. Lu, C., Krishna, R., Bernstein, M., Fei-Fei, L.: Visual relationship detection with language priors. In: European Conference on Computer Vision (2016)
18. Lu, Y., Rai, H., Chang, J., Knyazev, B., Yu, G., Shekhar, S., Taylor, G.W., Volkovs, M.: Context-aware scene graph generation with seq2seq transformers. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 15931–15941 (2021)
19. Ma, Y., Lv, W.: Identification of pneumonia in chest x-ray image based on transformer. Journal of Antennas and Propagation (2022)
20. Pino, P., Parra, D., Besa, C., Lagos, C.: 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. pp. 654–663. Springer (2021)
21. Rolínek, M., Swoboda, P., Zietlow, D., Paulus, A., Musil, V., Martius, G.: Deep graph matching via blackbox differentiation of combinatorial solvers. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVIII 16. pp. 407–424. Springer (2020)
22. Sharifzadeh, S., Baharlou, S.M., Tresp, V.: Classification by attention: Scene graph classification with prior knowledge. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 5025–5033 (2021)
23. Shin, H.C., Roberts, K., Lu, L., Demner-Fushman, D., Yao, J., Summers, R.M.: 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. pp. 2497–2506 (2016)
24. Shit, S., Koner, R., Wittmann, B., Paetzold, J., Ezhov, I., Li, H., Pan, J., Sharifzadeh, S., Kaissis, G., Tresp, V., et al.: Relationformer: A unified framework for image-to-graph generation. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVII. pp. 422–439. Springer (2022)

25. Wang, X., Peng, Y., Lu, L., Lu, Z., Summers, R.M.: 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. pp. 9049–9058 (2018)
26. Wills, P., Meyer, F.G.: Metrics for graph comparison: a practitioner's guide. Plos one **15**(2), e0228728 (2020)
27. Wu, J.T., Agu, N.N., Lourentzou, I., Sharma, A., Paguio, J.A., Yao, J.S., Dee, E.C., Mitchell, W., Kashyap, S., Giovannini, A., et al.: Chest imagenome dataset for clinical reasoning. arXiv preprint arXiv:2108.00316 (2021)
28. Yu, F., Endo, M., Krishnan, R., Pan, I., Tsai, A., Reis, E.P., Fonseca, E.K.U.N., Lee, H.M.H., Abad, Z.S.H., Ng, A.Y., et al.: Evaluating progress in automatic chest x-ray radiology report generation. medRxiv pp. 2022–08 (2022)
29. Yun, S., Jeong, M., Kim, R., Kang, J., Kim, H.J.: Graph transformer networks. Advances in neural information processing systems **32** (2019)

Supplementary material

Table 4: **Sample entity mappings.**

|	
clear	clear, lucencies, lucency
consolidation	consolidation, consolidations, consolidative
congested	congested, congestion, engorged
sharp	sharp, sharply
prominent	prominent, prominence
nodule	nodule, nodules, nodular
pneumonic	pneumonic, pneumonia
calcification	calcification, calcified, calcifications
tortuous	tortuous, tortuosity
atelectasis	atelectasis, atelectatic, atelectases
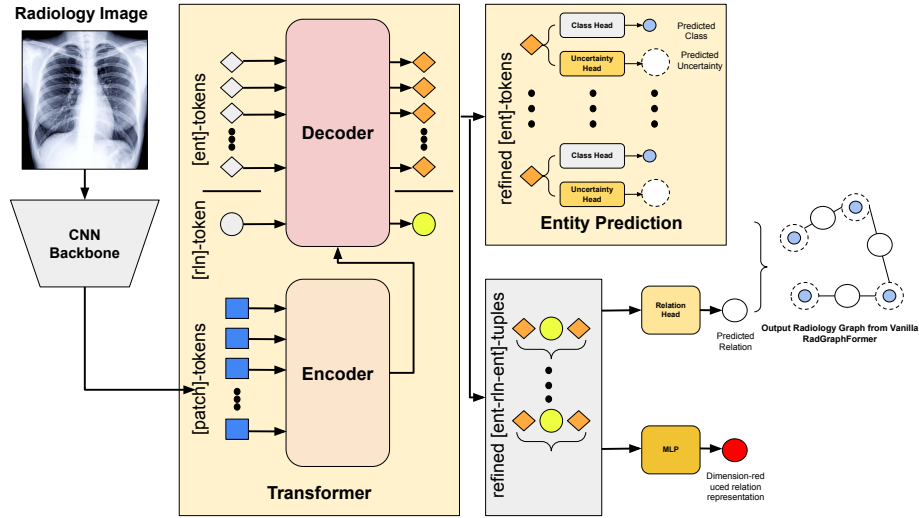
Fig. 3: **Vanilla RadGraphFormer architecture.** Vanilla RadGraphFormer passes the concatenation of a pair of [ent]-tokens and a shared [rln]-token (concatenated edge feature) to a relation head to predict the relation type. For constructing the initial embedding graph for prior knowledge integration, an MLP is used to reduce the concatenated edge feature to make both nodes and edges dimension identical.
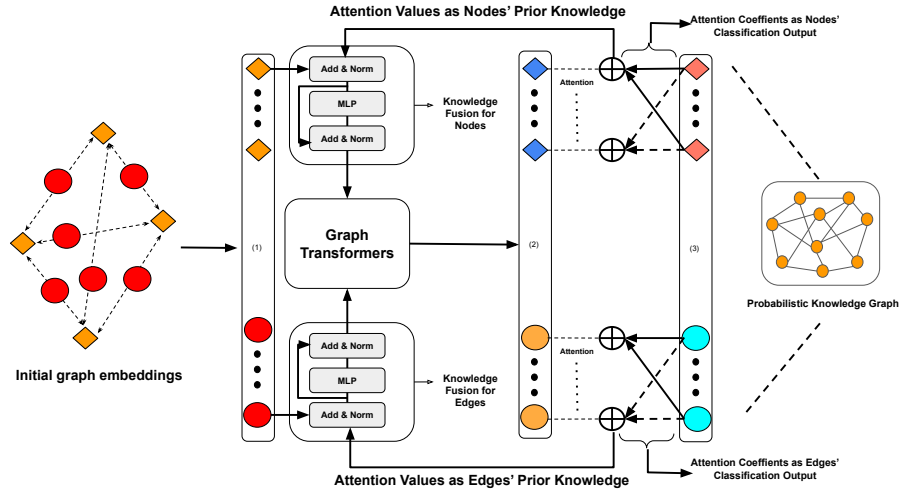

Fig. 4: **Prior knowledge integration process.** (1) Initial representations from vanilla RadGraphFormer; (2) Propagated features by graph transformers; (3) Randomly initialized schema representations.