
Inspecting state of the art performance and NLP metrics in image-based medical report generation

Pablo Pino ¹, Denis Parra ^{1,2,3}, Pablo Messina ¹, Cecilia Besa ^{2,3}, Sergio Uribe ^{2,3}

¹ Department of Computer Science, Pontificia Universidad Católica de Chile,

² School of Medicine, Pontificia Universidad Católica de Chile,

³ Millennium Nucleus in Cardiovascular Magnetic Resonance, ANID,
{pdpino, pamessina, cbesa, suribe}@uc.cl, dparra@ing.puc.cl

Abstract

Several deep learning architectures have been proposed over the last years to deal with the task of generating a written report given an imaging exam as input. Most works evaluate the generated reports using standard Natural Language Processing (NLP) metrics (e.g. BLEU, ROUGE), reporting significant progress. This article contrast this progress by comparing state of the art (SOTA) models against weak baselines. We show that simple and even naive approaches yield near SOTA performance on most traditional NLP metrics. We conclude that evaluation methods in this task should be further studied towards correctly measuring clinical accuracy, involving physicians to contribute to this end.

1 Introduction

Writing reports from medical images is a everyday labor for radiologists, like the given example of a chest X-rays shown in Figure 1. This task has already been addressed with deep learning approaches [e.g. 7, 9, 12, 14, 15], typically with encoder-decoder architectures employing a CNN to extract image features and an RNN for text generation. In most works, the authors evaluate the model performance using Natural Language Processing (NLP) metrics, such as BLEU [11], ROUGE-L [8] and CIDEr-D [13]. However, it has been discussed that these metrics cannot measure correctness in the medical domain [1, 9, 17], and that a small improvement in them may not implicate an actual quality improvement [10]. Some works have attempted to design a more appropriate evaluation method, for example, Liu et al. [9] applied the Chexpert Labeler [6] over the original and generated reports to classify 14 diseases, and used classification metrics, such as accuracy and ROC-AUC. In this work, we further inspect some NLP metrics by benchmarking different models against SOTA using the IU X-ray dataset [2], showing that weak baselines achieve comparable scores.



Original report (extract):

The cardiac silhouette **is enlarged**. The lungs are **hyperexpanded with flattening of the bilateral hemidiaphragms**. No pneumothorax or pleural effusion.

Possible constant report:

The cardiac silhouette **is normal in size and configuration**. The lungs are **clear**. No pneumothorax or pleural effusion.

Figure 1: Example from the IU X-ray dataset [2]: frontal chest X-ray and an extract of the radiologist’s report. The third column shows a possible constant report, which obtains high BLEU and ROUGE-L scores of 0.56 and 0.676, but describes incorrect clinical facts in two of its sentences.

2 Methods and Results

We compared weak baselines and CNN-LSTM architectures against SOTA [9, 14, 7, 15].

Weak baselines. Five baselines are implemented. (1) *Random*: selects a report randomly from the training set. (2) *Constant*: always outputs the same report, which we built with some of the most common sentences describing no findings (see example in Figure 1). (3) *Nearest-neighbor*: looks for the most similar image in the training set and output its report. (4) *Top-sentences-N* / (5) *Top-words-N*: selects the N most common sentences/words of the training set and returns them in random order.

CNN-LSTM. A CNN extracts features from the input image, which are fed to an LSTM that generates the report word by word. Each input word is passed through an embedding layer. Two variants are implemented, (1) *CNN-LSTM*: the image features are fed to a fully connected layer to reduce its size and to initialize the LSTM hidden state. (2) *CNN-LSTM-att*: same as (1), plus using an attention network over the spatial dimension of the image features at each step, and concatenating the attended features with the word embedding input. The attention network consists of two fully connected layers with a *tanh* activation, following the method used by Tian et al. [12]. We tested several CNNs: Mobilenet-v2 [4], Resnet-50 [3], Densenet-121 [5]; the latter performed the best.

2.1 Results

Table 1 shows the each model’s performance in terms of NLP metrics, besides accuracy and ROC-AUC calculated after applying the Chexpert Labeler to the reports (both macro-averaged across 14 diseases). All metrics range from 0 (worst) to 1 (best), except for CIDEr-D, which reaches a maximum of 10. The top three results per column are in bold, showing weak baselines achieve scores near to or better than SOTA models, considering BLEU and ROUGE-L. However, Chexpert metrics show that models are not detecting diseases correctly, as the baseline AUC for a binary classification task is 0.5. Notice that because most sentences in the reports describe no abnormalities (about 76% of the sentences in the dataset), simple baselines can achieve high accuracy. Also, CIDEr-D shows a strong difference between two SOTA models and naive baselines, though they obtain similar accuracy. These facts suggest NLP metrics may not be suitable for evaluating the generated reports, and further research should determine how to assess this task automatically.

Table 1: Model performance in the IU X-ray dataset [2]. Results are taken from each paper, except for TieNet which are taken from Liu et al. [9]. B: BLEU, R-L: ROUGE-L, C-D: CIDEr-D.

| Model | B-1 | B-2 | B-3 | B-4 | B | R-L | C-D | Acc | AUC |
|-------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| Liu et al. [9] | 0.369 | 0.246 | 0.171 | 0.115 | 0.225 | 0.359 | 1.490 | 0.916 | - |
| TieNet [14] | 0.330 | 0.194 | 0.124 | 0.081 | 0.182 | 0.311 | 1.335 | 0.902 | - |
| KERP [7] | 0.482 | 0.325 | 0.226 | 0.162 | 0.298 | 0.339 | 0.280 | - | - |
| Xue et al. [15] | 0.477 | 0.332 | 0.243 | 0.189 | 0.310 | 0.380 | - | - | - |
| Constant | 0.455 | 0.312 | 0.223 | 0.165 | 0.289 | 0.357 | 0.293 | 0.915 | 0.500 |
| Random | 0.362 | 0.197 | 0.117 | 0.075 | 0.188 | 0.264 | 0.112 | 0.894 | 0.508 |
| Nearest-neighbor | 0.383 | 0.220 | 0.142 | 0.100 | 0.211 | 0.288 | 0.230 | 0.903 | 0.518 |
| Top-sentences-100 | 0.347 | 0.211 | 0.138 | 0.096 | 0.198 | 0.281 | 0.166 | 0.911 | 0.498 |
| Top-words-50 | 0.375 | 0.102 | 0.019 | 0.000 | 0.124 | 0.224 | 0.075 | 0.835 | 0.509 |
| CNN-LSTM | 0.379 | 0.239 | 0.164 | 0.117 | 0.225 | 0.338 | 0.284 | 0.912 | 0.505 |
| CNN-LSTM-att | 0.361 | 0.226 | 0.152 | 0.106 | 0.211 | 0.314 | 0.187 | 0.918 | 0.508 |

3 Conclusions and future work

Given the task of image-based medical report generation, we compared multiple models to the SOTA, using NLP metrics and the Chexpert Labeler. We found that some weak baselines performed almost as well as some SOTA models in terms of BLEU and ROUGE-L, suggesting that either these models do not perform so well in a clinical sense, or that these NLP metrics are not able to differentiate automatic methods. Furthermore, we can raise multiple research questions, are SOTA models actually performing better than weak baselines in a clinical sense? Do NLP metrics correlate with medical judgement for this task? How to directly measure medical correctness in a generated report? In future work, we will address these questions by asking radiologists to assess generated reports, and we will test the metric MIRQI, which was recently proposed by Zhang et al. [16] and is designed to measure clinical correctness in the written reports.

Broader Impact

This research attempts to generate further discussion on how to measure progress on medical report generation from medical images. Despite the lack of scalable resources to judge generated reports' quality (expert physicians), we believe the gold standard should consider expert human judgment to assess these methods. Several methods use traditional metrics such as ROUGE or BLEU, but since they have been contested in tasks such as Machine Translation [10], we suggest to further investigate how appropriate it is to use them for assessing medical report generation. We propose moving forward to metrics actually measuring medical accuracy rather than traditional Natural Language Generation (NLG). We expect to positively impact this area of research, by providing new metrics to further impact the progress of the application of NLP to medical reporting.

Acknowledgments and Disclosure of Funding

This work has been supported by the Millennium Institute for Foundational Research on Data (IMFD) and by the Chilean research agency ANID, with the FONDECYT grant 1191791, and with the Millennium Science Initiative Program, NCN17_129.

References

- [1] W. Boag et al. "Baselines for Chest X-Ray Report Generation". In: *Proc of the Machine Learning for Health NeurIPS Workshop*. 2020, pp. 126–140.
- [2] D. Demner-Fushman et al. "Preparing a collection of radiology examinations for distribution and retrieval". In: *Journal of the American Medical Informatics Assoc.* (2015), pp. 304–310.
- [3] K. He et al. "Deep residual learning for image recognition". In: *Proc of the IEEE Conf. on CVPR*. 2016, pp. 770–778.
- [4] A. G. Howard et al. "Mobilenets: Efficient convolutional neural networks for mobile vision applications". In: *arXiv:1704.04861* (2017).
- [5] G. Huang et al. "Densely connected convolutional networks". In: *Proc of the IEEE Conf. on CVPR*. 2017, pp. 4700–4708.
- [6] J. Irvin et al. "Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison". In: *Proc of the AAAI Conf. on Artificial Intelligence*. 2019, pp. 590–597.
- [7] C. Y. Li et al. "Knowledge-Driven Encode, Retrieve, Paraphrase for Medical Image Report Generation". In: *Proc of the AAAI Conf. on Artificial Intelligence*. 2019, pp. 6666–6673.
- [8] C.-Y. Lin. "ROUGE: A Package for Automatic Evaluation of Summaries". In: *Text Summarization Branches Out*. 2004, pp. 74–81.
- [9] G. Liu et al. "Clinically Accurate Chest X-Ray Report Generation". In: *Machine Learning for Healthcare Conf*. 2019, pp. 249–269.
- [10] N. Mathur, T. Baldwin, and T. Cohn. "Tangled up in BLEU: Reevaluating the Evaluation of Automatic Machine Translation Evaluation Metrics". In: *Proc of the 58th Annual Meeting of the ACL*. 2020, pp. 4984–4997.
- [11] K. Papineni et al. "Bleu: a Method for Automatic Evaluation of Machine Translation". In: *Proc of the 40th Annual Meeting of the ACL*. 2002, pp. 311–318.
- [12] J. Tian et al. "A Diagnostic Report Generator from CT Volumes on Liver Tumor with Semi-supervised Attention Mechanism". In: *MICCAI*. 2018, pp. 702–710.
- [13] R. Vedantam, C. Lawrence Zitnick, and D. Parikh. "Cider: Consensus-based image description evaluation". In: *Proc of the IEEE Conf. on CVPR*. 2015, pp. 4566–4575.
- [14] X. Wang et al. "Tienet: Text-image embedding network for common thorax disease classification and reporting in chest x-rays". In: *Proc of the IEEE Conf. on CVPR*. 2018.
- [15] Y. Xue and X. Huang. "Improved Disease Classification in Chest X-Rays with Transferred Features from Report Generation". In: *IPMI*. 2019, pp. 125–138.
- [16] Y. Zhang et al. "When Radiology Report Generation Meets Knowledge Graph". In: *arXiv:2002.08277* (2020).
- [17] Y. Zhang et al. "Optimizing the Factual Correctness of a Summary: A Study of Summarizing Radiology Reports". In: *Proc of the 58th Annual Meeting of the ACL*. 2020, pp. 5108–5120.