Unsupervised Multimodal Representation Learning across Medical Images and Reports

Tzu-Ming Harry Hsu, Wei-Hung Weng, Willie Boag, Matthew McDermott, Peter Szolovits MIT CSAIL {stmharry,ckbjimmy,wboag,mmd,psz}@mit.edu

Abstract

Joint embeddings between medical imaging modalities and associated radiology reports have the potential to offer significant benefits to the clinical community, ranging from cross-domain retrieval to conditional generation of reports to the broader goals of multimodal representation learning. In this work, we establish baseline joint embedding results measured via both local and global retrieval methods on the soon to be released MIMIC-CXR dataset consisting of both chest X-ray images and the associated radiology reports. We examine both supervised and unsupervised methods on this task and show that for document retrieval tasks with the learned representations, only a limited amount of supervision is needed to yield results comparable to those of fully-supervised methods.

1 Introduction

Medical imaging is one of the most compelling domains for the immediate application of artificial intelligence tools. Recent years have seen not only tremendous academic advancements [8, 12, 22] but additionally a breadth of applied tools [7, 16, 18, 29].

There has been some emerging attention on joint processing of medical images and radiological free-text reports. [31] used the public NIH Chest X-ray 14 dataset [30] linked with the non-public associated reports to both improve disease classification performance and for automatic report generation. [9] attempted to generate radiology reports while [27] generated disease/location/severity annotations. [17] generated notes, including radiology reports for the Medical Information Mart for Intensive Care (MIMIC) dataset using non-image modalities such as demographics, previous notes, labs, and medications. These works used annotations from either machines [30] or humans. However, with a huge influx of imaging data beyond human capacity, parallel records from both imaging and text are not always readily available. We thus would like to bring up the question of whether we can take advantage of unannotated but massive imaging datasets and learn from the underlying distribution of these images.

One natural area that remains unexplored is representation learning across images and reports. The idea of representation learning in a joint embedding space can be realized in multiple ways. Some [4, 20] explored statistical and metrical relevance across domains, and some [10] realized it as an adversarially determined domain-agnostic latent space. [19, 26] both used a the latent space for style transfer, in language sentiment and music style, respectively. [23] learned joint spaces of images and their captions, which [24] later used for caption-driven image generation. [6] and [11] also used similar ideas to perform both supervised and unsupervised word-to-word translation tasks.

A recent dataset, MIMIC- Chest X-ray (MIMIC-CXR), carries paired records of X-ray images and radiology reports, and the imaging modality has been explored in [25]. In this work, we explore

Machine Learning for Health (ML4H) Workshop at NeurIPS 2018.

¹Soon to be publicly released.



Figure 1: The overall experimental pipeline. EA: embedding alignment; Adv: adversarial training.

both the text and image modalities with joint embedding spaces under a spectrum of supervised and unsupervised methods. In particular, we make the following contributions:

- 1. We establish baseline results and evaluation methods for jointly embedding radiological images and reports via retrieval and distance metrics.
- 2. We profile the impact of supervision level on the quality of representation learning in joint embedding spaces.
- 3. We characterize the influence of using different sections from the report on representation learning.

2 Methodology

2.1 Data

All experiments in this work used the MIMIC-CXR dataset. MIMIC-CXR consists of 473,057 chest X-ray images and 206,563 reports from 63,478 patients. Of these images, 240,780 are of anteroposterior (AP) views, which we focus on in this work. Further, we eliminate all duplicated radiograph images with adjusted brightness or contrast? leaving a total of 95,242/87,353 images/reports, which we subdivide into a train set of 75,147/69,171 and a test set of 19,825/18,182 images/reports, with no overlap of patients between the two. Radiological reports are parsed into sections and we use either the *impression* or the *findings* sections.

For evaluation, we aggregate a list of unique International Classification of Diseases (ICD-9) codes from all patient admissions and ask a clinician to pick out a subset of codes that are related to thoracic diseases. Records with ICD-9 codes in the subset are then extracted, including 3,549 images from 380 patients. This population serves as a disease-related evaluation for retrieval algorithms. Note that this disease information is never provided during training in any setting.

2.2 Methods

Our overall experimental flow follows Figure []. Notes are featurized via (1) term frequencyinverse document frequency (TF-IDF) over bi-grams, (2) pre-trained GloVe *word* embeddings [21] averaged across the selected section of the report, (3) *sentence* embeddings, or (4) *paragraph* embeddings. In (3) and (4), we first perform sentence/paragraph splitting, and then fine-tune a deep averaging network (DAN) encoder []] [3] [13] with the corpus. Embeddings are finally averaged across sentences/paragraphs. The DAN encoder is pretrained on a variety of data sources and tasks and fine-tuned on the context of report sections.

Images are resized to 256×256 , then featurized to the last bottleneck layer of a pretrained DenseNet-121 model [22]. PCA is applied onto the 1024-dimension raw image features to obtain 64-dimension features ³ Text features are projected into the 64-dimension image feature space. We use several methods regarding different objectives.

Embedding Alignment (EA) Here, we find a linear transformation between two sets of matched points $\mathbf{X} \in \mathbb{R}^{d_X \times n}$ and $\mathbf{Y} \in \mathbb{R}^{d_Y \times n}$ by minimizing $\mathcal{L}_{\text{EA}}(\mathbf{X}, \mathbf{Y}) = \|\mathbf{W}^\top \mathbf{X} - \mathbf{Y}\|_F^2$.

²Commonly produced for clinical needs

³96.9% variance explained

Text Feature	Method	Similarity	$MRR(\times 10^{-3})$		nDCG@1		nDCG@10		nDCG@100	
			$T\toI$	$I\toT$	$T\toI$	$I\toT$	$T\toI$	$I\toT$	$T \rightarrow I$	$I\toT$
	chance		0.50	0.50	0.103	0.103	0.103	0.103	0.103	0.103
bi-gram	EA	0.613 .000	$7.33_{.04}$	11.65 .07	$0.147_{.001}$	$0.162_{.001}$	0.148 .000	0.159 .000	0.225 .000	0.231 .000
word	EA	$0.542_{.000}$	$2.00_{.01}$	$4.52_{.02}$	$0.096_{.002}$	$0.128_{.001}$	$0.116_{.000}$	$0.130_{.000}$	$0.202_{.000}$	$0.205_{.000}$
sentence	EA	$0.465_{.000}$	$1.08_{.00}$	$2.74_{.02}$	$0.073_{.001}$	$0.101_{.000}$	0.100_{-000}	$0.111_{.000}$	$0.189_{.000}$	$0.177_{.000}$
paragraph	EA	$0.505_{.000}$	$1.57_{.01}$	$2.53_{.01}$	$0.082_{.001}$	$0.134_{.000}$	$0.107_{.000}$	$0.124_{.000}$	$0.195_{.000}$	$0.196_{.000}$
bi-gram	Adv	$0.218_{.073}$	$0.77_{.23}$	$0.85_{.33}$	$0.095_{.006}$	$0.090_{.003}$	$0.101_{.004}$	$0.098_{.003}$	$0.171_{.005}$	$0.166_{.004}$
bi-gram	Adv + Proc	$0.221_{.074}$	$0.77_{.24}$	$0.87_{.32}$	$0.094_{.006}$	$0.091_{.004}$	$0.102_{.004}$	$0.099_{.002}$	$0.171_{.005}$	$0.166_{.004}$
word	Adv	$0.268_{.016}$	$0.65_{.12}$	$0.54_{.12}$	$0.096_{.006}$	$0.091_{.003}$	$0.105_{.004}$	$0.099_{.003}$	$0.176_{.003}$	$0.165_{.004}$
word	Adv + Proc	$0.269_{.013}$	$0.64_{.11}$	$0.57_{.07}$	$0.098_{.006}$	$0.092_{.002}$	$0.107_{.005}$	$0.099_{.003}$	$0.179_{.003}$	$0.165_{.004}$
sentence	Adv	$0.265_{.010}$	$0.64_{.08}$	$1.07_{.24}$	$0.095_{.007}$	$0.094_{.002}$	$0.103_{.006}$	$0.100_{.001}$	$0.176_{.006}$	$0.167_{.001}$
sentence	Adv + Proc	$0.266_{.012}$	$0.68_{.10}$	$1.07_{.21}$	$0.096_{.005}$	$0.094_{.004}$	$0.105_{.006}$	$0.100_{.002}$	$0.178_{.005}$	$0.166_{.002}$
paragraph	Adv	$0.045_{.136}$	$0.69_{.03}$	$0.70_{.04}$	$0.062_{.025}$	$0.123_{.029}$	$0.082_{.015}$	$0.118_{.017}$	$0.163_{.013}$	$0.169_{.003}$
paragraph	Adv + Proc	$0.225_{.061}$	$1.15_{.60}$	$0.77_{.21}$	$0.093_{.057}$	$0.092_{.011}$	$0.090_{.034}$	$0.103_{.008}$	$0.163_{.023}$	$0.166_{.005}$

Table 1: Comparison among supervised (upper) and unsupervised (lower) methods. Subscripts show the half width of 95% confidence intervals. **Bold** denotes the best performance in each group. *Chance* is the expected value if we randomly yield retrievals. Higher is better for all metrics.

Adversarial Domain Adaption (Adv) Adversarial training pits a *discriminator*, D, implemented as a 2-layer (hidden size 256) neural network using scaled exponential linear units (SELUs) [15], against a projection matrix \mathbf{W} , as the *generator*. D is trained to classify points in the joint space according to source modality, and \mathbf{W} is trained adversarially to fool D. Alternatively, D minimizes $\mathcal{L}_{Adv}^{D}(\mathbf{X}, \mathbf{Y}) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim p(\mathbf{X}, \mathbf{Y})} \left[-\log D(\mathbf{W}^{\top} \mathbf{x}) - \log (1 - D(\mathbf{y})) \right]$ when \mathbf{W} minimizes $\mathcal{L}_{Adv}^{\mathbf{W}}(\mathbf{X}, \mathbf{Y}) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim p(\mathbf{X}, \mathbf{Y})} \left[-\log \left(1 - D(\mathbf{W}^{\top} \mathbf{x}) \right) \right]$.

Procrustes Refinement (Adv + Proc) On top of adversarial training, we also use an unsupervised Procrustes induced refinement as in [6].

Semi-Supervised We also assess how much supervision is necessary to ensure strong performance on these modalities by randomly subsampling our data into supervised and unsupervised samples. We then combine the *embedding alignment* objective and *adversarial training* objective functions as $\mathcal{L} = \mathcal{L}_{EA}(\mathbf{X}, \mathbf{Y}) + \lambda \mathcal{L}_{Adv}(\mathbf{X}, \mathbf{Y})$ and train simultaneously as we vary the fraction trained. Preliminary experiments suggests $\lambda = 0.1$.

Orthogonal Regularization [6, [28, [32]] all showed that imposing orthonormality on linear projections leads to better performance and stability in training. However, [2] suggested orthogonality (i.e., not constraining the norms) can perform better as a regularization. Thus on top of the objectives, we add $\mathcal{R}_{ortho} = \beta \| \mathbf{W}^\top \mathbf{W} \odot (\mathbf{e}\mathbf{e}^\top - \mathbf{I}) \|_F^2$, where \odot denotes element-wise product and \mathbf{e} denotes a column vector of all ones. Scanning through a range shows $\beta = 0.01$ yields good performance.

2.3 Evaluation

We evaluate via cross domain retrieval in the test set Q: querying in the joint embedding space for closest neighboring images using a report, $T \rightarrow I$, or vice-versa, $I \rightarrow T$. For direct pairings, we compute the *cosine similarity*, and MRR $= \frac{1}{|Q|} \sum_{q \in Q} \frac{1}{rank_q}$ where rank_q is the rank of the first true pair for q (e.g., the first paired image or text corresponding to the query q) in the retrieval list. For thoracic disease induced pairings, we first define the relevance $\operatorname{rel}_{pq} \in [0, 1]$ between two entries p and q as the intersection-over-union of their respective set of ICD-9 codes. Then we calculate the normalized discounted cumulative gain [I4] nDCG@k $= \frac{1}{|Q|} \sum_{q \in Q} \frac{1}{\operatorname{IDCG}_q} \sum_{p=1}^k \frac{2^{\operatorname{rel}pq}-1}{\log_2(p+1)}$, where IDCG_q denotes the ideal DCG value for q using a perfect retrieval algorithm. All experiments are repeated with random initial seeds for at least 5 times. Means and 95% confidence intervals are reported in the following section.

3 Results

Retrieval with/without Supervision Table 1 compares four types of text features and supervised/unsupervised methods. We find that unsupervised methods can achieve comparable results on disease-related retrieval tasks on a large scale (nDCG@100) without the need for labeling the chest X-ray images. Experiments show uni-, bi-, and tri-grams yield very similar results and we

only include bi-gram in the table. Additionally, we find that the high-level *sentence* and *paragraph* embeddings approach underperformed the bi-gram text representation. Although having generalizability [3], sentence and paragraph embeddings learned from the supervised multi-task pre-trained model may not be able to represent the domain-specific radiological reports well due to the lack of medical domain tasks in the pre-training process. Unsupervised procrustes refinement is occasionally, but not universally helpful. Note that MRR is comparatively small since reports are in general highly similar for radiographs with the same disease types.

The Impact of Supervision Fraction We define the *supervision fraction* as the fraction of pairing information provided in the training set. Note the ICD-9 codes are not provided for training even in the fully supervised setting. Figure 2 shows our evaluation metrics for models trained using bi-gram text features and the semi-supervised learning objective for various supervision fractions. A minimal supervision as low as 0.1% provided can drastically improve the alignment quality, especially in terms of cosine similarity and nDCG. More annotations further improve the performance measures, but one would almost require exponentially many data points in exchange for a linear increase. That implies the possibility of concatenating a well-annotated dataset and a large but unannotated dataset for a substantial performance boost.



Figure 2: Performance measures of retrieval tasks at k retrieved items as a function of the supervision fraction. Higher is better. Note the *x*-axis is in log scale. Unsupervised is on the left, increasingly supervised to the right. Dashed lines indicate the performance by chance. Vertical bars indicate the 95% confidence interval, and some are too narrow to be visible.

Using Different Sections of the Report We investigate the effectiveness of using different sections for the embedding alignment task. All models in Figure **3** run with a supervision fraction of 1%. The models trained on the *findings* section outperformed the models trained on the *impression* section using cosine similarity and MRR. This makes sense from a clinical perspective since the radiologists usually only describe image patterns in the *findings* section and thus they would be aligned well. On the other hand, they make radiological-clinical integrated interpretations in the *impression* section, which means that the both the image-uncorrelated clinical history and findings were mentioned in the *impression* section. Since nDCG is calculated using ICD-9 codes, which carry disease-related information, it naturally aligns with the purpose of writing an *impression* section. This may explain why the models trained on *impression* section worked better for nDCG.



Figure 3: Different metrics for retrieval on either the *impression* or *findings* section using four types of features. 95% confidence intervals are indicated on the bars.

4 Conclusion

MIMIC-CXR will soon be the largest publicly available imaging dataset consisting of both medical images and paired radiological reports, promising myriad applications that can make use of both modalities together. We establish baseline results using supervised and unsupervised joint embedding methods along with local (direct pairs) and global (ICD-9 code groupings) retrieval evaluation metrics. Results show a possibility of incorporating more unsupervised data into training for minimal-effort performance increase. A further study of joint embeddings between these modalities may enable significant applications, such as text/image generation or the incorporation of other EMR modalities.

References

- [1] Steven Bird and Edward Loper. NLTK: the natural language toolkit. In *Proceedings of the ACL 2004 on Interactive poster and demonstration sessions*, page 31. Association for Computational Linguistics, 2004.
- [2] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- [3] Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. Universal sentence encoder. arXiv preprint arXiv:1803.11175, 2018.
- [4] Wei-Yu Chen, Tzu-Ming Harry Hsu, Yao-Hung Hubert Tsai, Yu-Chiang Frank Wang, and Ming-Syan Chen. Transfer neural trees for heterogeneous domain adaptation. In *European Conference on Computer Vision*, pages 399–414. Springer, 2016.
- [5] Yu-An Chung, Wei-Hung Weng, Schrasing Tong, and James Glass. Unsupervised cross-modal alignment of speech and text embedding spaces. arXiv preprint arXiv:1805.07467, 2018.
- [6] Alexis Conneau, Guillaume Lample, Marc'Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. Word translation without parallel data. *arXiv preprint arXiv:1710.04087*, 2017.
- [7] EnvoyAI. EnvoyAI launches with 35 algorithms contributed by 14 newly-contracted artificial intelligence development partners. *EnvoyAI Blog*, 11 2017.
- [8] Andre Esteva, Brett Kuprel, Roberto A Novoa, Justin Ko, Susan M Swetter, Helen M Blau, and Sebastian Thrun. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639):115, 2017.
- [9] William Gale, Luke Oakden-Rayner, Gustavo Carneiro, Andrew P Bradley, and Lyle J Palmer. Producing radiologist-quality reports for interpretable artificial intelligence. *arXiv preprint arXiv:1806.00340*, 2018.
- [10] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016.
- [11] Edouard Grave, Armand Joulin, and Quentin Berthet. Unsupervised alignment of embeddings with wasserstein procrustes. *arXiv preprint arXiv:1805.11222*, 2018.
- [12] Varun Gulshan, Lily Peng, Marc Coram, Martin C Stumpe, Derek Wu, Arunachalam Narayanaswamy, Subhashini Venugopalan, Kasumi Widner, Tom Madams, Jorge Cuadros, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA*, 316(22):2402–2410, 2016.
- [13] Mohit Iyyer, Varun Manjunatha, Jordan Boyd-Graber, and Hal Daumé III. Deep unordered composition rivals syntactic methods for text classification. In Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), volume 1, pages 1681–1691, 2015.
- [14] Kalervo Järvelin and Jaana Kekäläinen. Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems (TOIS)*, 20(4):422–446, 2002.
- [15] Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. Self-normalizing neural networks. In Advances in Neural Information Processing Systems, pages 971–980, 2017.
- [16] Jeff Lagasse. FDA approves first AI tool for detecting retinopathy, NIH shows machine learning success in imaging | Healthcare Finance News. *Healthcare Finance*, 4 2018.
- [17] Peter J Liu. Learning to write notes in electronic health records. *arXiv preprint arXiv:1808.02622*, 2018.
- [18] Bernard Marr. First FDA approval for clinical cloud-based deep learning in healthcare. *Forbes*, 2017.
- [19] Noam Mor, Lior Wolf, Adam Polyak, and Yaniv Taigman. A universal music translation network. *arXiv preprint arXiv:1805.07848*, 2018.
- [20] Sinno Jialin Pan, Ivor W Tsang, James T Kwok, and Qiang Yang. Domain adaptation via transfer component analysis. *IEEE Transactions on Neural Networks*, 22(2):199–210, 2011.

- [21] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP), pages 1532–1543, 2014.
- [22] Pranav Rajpurkar, Jeremy Irvin, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, Daisy Ding, Aarti Bagul, Curtis Langlotz, Katie Shpanskaya, et al. CheXnet: Radiologist-level pneumonia detection on chest X-rays with deep learning. arXiv preprint arXiv:1711.05225, 2017.
- [23] Scott Reed, Zeynep Akata, Honglak Lee, and Bernt Schiele. Learning deep representations of fine-grained visual descriptions. In *Proceedings of the IEEE Conference on Computer Vision* and Pattern Recognition, pages 49–58, 2016.
- [24] Scott Reed, Zeynep Akata, Xinchen Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. arXiv preprint arXiv:1605.05396, 2016.
- [25] Jonathan Rubin, Deepan Sanghavi, Claire Zhao, Kathy Lee, Ashequl Qadir, and Minnan Xu-Wilson. Large scale automated reading of frontal and lateral chest x-rays using dual convolutional neural networks. arXiv preprint arXiv:1804.07839, 2018.
- [26] Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. Style transfer from nonparallel text by cross-alignment. In *Advances in Neural Information Processing Systems*, pages 6830–6841, 2017.
- [27] Hoo-Chang Shin, Kirk Roberts, Le Lu, Dina Demner-Fushman, Jianhua Yao, and Ronald M Summers. Learning to read chest X-rays: Recurrent neural cascade model for automated image annotation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2497–2506, 2016.
- [28] Samuel L Smith, David HP Turban, Steven Hamblin, and Nils Y Hammerla. Offline bilingual word vectors, orthogonal transformations and the inverted softmax. *arXiv preprint arXiv:1702.03859*, 2017.
- [29] Michael Walter. Densitas gains FDA clearance for machine learning software that assesses breast density. *Radiology Business*, 4 2018.
- [30] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald M Summers. ChestX-ray8: Hospital-scale chest X-ray database and benchmarks on weaklysupervised classification and localization of common thorax diseases. In *Computer Vision and Pattern Recognition (CVPR)*, 2017 IEEE Conference on, pages 3462–3471. IEEE, 2017.
- [31] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, and Ronald M Summers. 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*, pages 9049–9058, 2018.
- [32] Chao Xing, Dong Wang, Chao Liu, and Yiye Lin. Normalized word embedding and orthogonal transform for bilingual word translation. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1006–1011, 2015.