Assessing Ventilator-Associated Pneumonia (VAP) in the PICU
Project Proposal

Students: Suraj Shah

Mentors: Dr. Matias Unberath and RAIL (Radiology Artificial Intelligence Lab)

Clinical Collaborators: Drs. Dr. Jim Fackler, Jules Bergmann, Ferdinand Hui, Haris Sair, Paul Yi

Objective: Our objective is to determine if we can identify VAP risk early through analyzing time series data and other identification markers gathered in the PICU. There is a multidisciplinary team working on addressing this issue; we are specifically working with the radiology team to focus on X-ray image changes associated with VAP risk. Our goal is to develop an algorithm that can accurately predict if a patient has VAP by analyzing time-series images throughout a patient's addressed history.

Clinical Background:
Ventilator-associated Pneumonia (VAP), “specifically refers to pneumonia developing in a mechanically ventilated patient more than 48 hours after tracheal intubation.”[1] When under ventilation, patients are in a critical, life-sustaining ICU therapy; the body is at a fragile state and susceptible to diseases, including bacterial infections that are attributed to VAP. There are other risks involved as well, as further disease progression, volume overload, latrogenic infection, and ventilator injury. VAP has had such a deteriorating effect in the ICU ward that is now the leading cause of mortality among nosocomial infections and the leading cause of nosocomial morbidity[1]. Specifically, acquiring VAP increases the risk of morbidity by 30% for any given patient. This is compounded by the fact that 10-20% of ICU patients are diagnosed with VAP annually.
Ventilator-associated complications are also correlated with a much greater length of stay and time under ventilation. This leads to greater strain on the entire healthcare value chain, from the provider, insurer, and most importantly, the patient.

Motivation:
While there is a multi-disciplinary team working on addressing the issues arising from VAP in the ICU (including the PICU team working on identifying biomarkers and the ID team working on appropriate cultures and antibiotics), there hasn’t been a comprehensive study connecting the radiology component of monitoring patients with VAP and/or risk of VAP. This is mainly due to the clinical data collection process. The X-ray images that are collected occur over many different hospitals, at different orientations of patients, on different machines, and either at inspiry or expiry. There is not a standardized process for data collection, thus leading to an aggregation of thousands of X-ray images, but ones that are not easily ingestible into an algorithm that is able to readily classify the risk of VAP for a specific patient. For example, the largest chest X-ray dataset of adult images, MIMICS III, contains over 224,000 images of over 60,000 patients at various times during the monitoring phase in the ICU. Furthermore, when applying an algorithm to pediatric patients, we must understand that the patients grow much more quickly over time, thus any algorithm must account for the changes in image dimension. Thus, our project is to propose an algorithm that analyzes these images (from a significant database), and applies image transformation and normalization techniques, so that appropriate neural network package can be trained to accurately diagnose the occurrence of VAP for high-risk patients. This will lead to a renewed focus in the ICU on these high-risk patients, while avoiding unnecessary therapy for low-risk patients.

Technical Approach

The project workflow is split into the following main components:

1. Screening and Collecting Working Data to Build Database of X-Ray Images
2. Assembling an Input/Output Module for Analyzing Images within the Database
3. Developing Registration Techniques and Combining Into a Sophisticated Algorithm for Time-Series Data (if time permits)

The bulk of our time will be focused on the second component, testing models and determining the best-performing neural network to classify VAP occurrence on a static image.

Data Aggregation and Screening

First, we will be assembling a MySQL database of chest X-Ray images which will be hosted through MARCC (Maryland Advanced Research Computing Center). This database will comprise of data collected from publicly available datasets such as MIMICS III, Stanford, NIH, etc. We are focusing first on the publicly available datasets as we await IRB approval for the pediatric data from JHU. These datasets will be screened for abnormalities and outliers, as well as making sure we have sufficient time-series data for an acceptable number of patients, as we will not be able to make predictions from the time-series data without a large enough sample size.

Assembling Input/Output Module – Neural Net Classification

To first get a pulse for how our system will analyze chest X-ray images, we will need to create a static image predictor – namely being, choosing an apt neural network to correctly classify a patient's occurrence of VAP based off one image. Once we can train and identify the best performing neural net for static image prediction, we will be able to use the time-series images for a more accurate prognosis. Potential neural networks to consider include:

- VGG: developed in 2014, this network consists of 16 layers, and is appealing because of the uniform architecture. In terms of performance, it is one of the most preferred choices extracting features from images, which could be very relevant when pursuing relevant factors in the chest X-ray images. One concern with VGG is that has 138 million parameters, which could be a potential hurdle in terms of difficulty to handle.

- ResNet: also known as Residual Neural Network, this was developed in 2015 and focuses heavily on on “skip connections” (gated recurrent units used in RNNs) and batch normalization. The architecture consists of 152 layers but is much less parameter dependent, thus reducing the complexity involved.
- DenseNet: an extension of ResNet. DenseNet’s methodology proposes concatenating outputs from the previous layers instead of using the summation (ResNet merges previous layers with future leaders). The main goal here is to reduce the number of links between layers for more efficient and accurate processing.

Before ingesting the data into the training set pipeline, we will be analyzing each neural net and the strength of each (based off image properties). We will start by training the networks on a specified training dataset and subsequently testing it on specified training set. After determining a set accuracy threshold with our mentors and clinical collaborators, we can identify the best performing neural net. We will have to fine-tune the parameters of the neural net to best apply to our data, and perhaps split the data more rigorously (i.e. omitting multiple copies of the same patient) so that we are able to have the best performing model possible.

Developing Registration Techniques and Combining Into a Sophisticated Algorithm for Time-Series Data

In order to track clinically relevant changes over time, we need to first regularize and normalize each of the chest x-ray images. Because these images come from a variety of datasets from different hospitals, there would be a fair amount of noise and detection of clinically insignificant noise if we try to run neural networks on the raw data. We will instead preprocess using image transformation and 2d non-rigid registration methods. These transformations should be non-rigid as patients are oriented in 3d positions between films, however because we only have 2d images, the best way to regularize would be with 2d non rigid transformations. We will use and combine existing algorithms that track objects over a series of frames and transform chest x-ray images to the same ideal orientation for every patient, thereby optimizing the data for neural net training.

If time permits and we are able accurately perform transformation and registration techniques to the chest X-ray images, we will be able to perform classification with time-series data. Using the best performing neural net, we can analyze multiple image of a patient and predict (hopefully) with a higher accuracy, a patient’s occurrence and risk of VAP. This will be increasingly relevant to keep feature recognition consistent across images, as we will have a set of highly important features that will be used in our

prediction. One possible method of employing this technique is using Blazingly Fast Video Object Segmentation with Pixel-Wise Metric Learning. As shown here, this algorithm is able to detect relevant features of an object across multiple frames, which is especially for videos. However, since we will only be analyzing at most in tens of images for patients, we believe this algorithm will be even more robust in object identification across the images.

Deliverables

Documentation:

- Python/MATLAB source code
- Code documentation
- Database of X-Ray cohorts
- Report describing methods and achievements

Minimum

- A database of X-ray cohorts and segment them based off patient type, time series characteristics, and other important features
- Algorithm that can produce one image prediction

Expected

- Algorithm for working image alignment with subsequent classification

Maximum

- Sophisticated algorithm that handles arbitrary time series data for accurate prediction and monitoring

Dependencies

Dependencies	Solution	Expected Date	Effect if not Completed
Access to MIMICS Dataset	Complete Series of HIPAA Training Modules	Completed	Cannot start training of network
Access to Workbench with Significant GPU Processing Speed	Find appropriate computer in either Hackerman or Krieger	Completed	Unable to run neural network on datasets
Access to Google Cloud Credits	Request through Dr. Unberath	Completed	Unable to host the instance of the database
Access to image processing neural networks	Download through internet and/or email Dr. Unberath	Ongoing	Will have to run through other frameworks
IRB Approval for JHU Data	Submit through Dr. Unberath and Dr. Fackler	Ongoing	N/A
Access to Maryland Supercomputer (training and testing on entire database)	Request through Dr. Unberath	Ongoing	Will delay hosting of database and processing time

Schedule

Dates	Activity	Completion
Week of 2/15	Prepare for in-class presentation and finalize project	Yes
Week of 2/22	Work off of in-class presentation	Yes
Week of 3/1	Google cloud account set up with all images uploaded Shadow Clinicians Spreadsheet of datasets with important Written project proposal due	In Process
Week of 3/8	Start hosting database on MARCC	No
Week of 3/15	Create database	No
Week of 3/22	Spring break	No
Week of 3/29	Run scripts to pull and clean data; start implementing training of pre-existing models	No
Week of 4/5	Continue to train pre-existing models and fine-tune parameters	No
Week of 4/12	Select best-performing model and start on feature recognition	No
Week of 4/19	Implement feature registration (Blazingly Fast Image Segmentation) to track clinically relevant features	No
Week of 4/26	Testing and poster/paper work (combine algorithms if time permits)	No
Week of 5/3	Poster presentation and final paper deliverables	No

Management Plan

- Thursday weekly meetings with Dr. Unberath & RAIL (Radiology Artificial Intelligence Lab)
- Communication with Dr. Jules Bergmann and Dr. Jim Fackler (in coordination with Dr. Unberath)
    - Email and text
- Code on private github
- Data stored in Google cloud account

Bibliography and Reading List

1. Kollef, M. H., Dr. (2005). WHAT IS VENTILATOR-ASSOCIATED PNEUMONIA AND WHY IS IT IMPORTANT? Respiratory Care, 50(6), 714-724. Retrieved February 26, 2019
2. Chen, Y., Pont-Tuset, J., Montes, A. and Van Gool, L. (2019). *Blazingly Fast Video Object Segmentation with Pixel-Wise Metric Learning*. [online] arXiv.org. Available at: https://arxiv.org/abs/1804.03131 [Accessed 18 Feb. 2019].
3. Yin Y, Hoffman EA, Ding K, Reinhardt JM, Lin CL. A cubic B-spline-based hybrid registration of lung CT images for a dynamic airway geometric model with large deformation. *Phys Med Biol*. 2010;56(1):203-18.
4. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770-778.
5. He, K., Zhang, X., Ren, S. and Sun, J. (2019). *Identity Mappings in Deep Residual Networks*. [online] arXiv.org. Available at: https://arxiv.org/abs/1603.05027 [Accessed 18 Feb. 2019].
6. Hatami, N., Gavet, Y. and Debayle, J. (2019). *Classification of Time-Series Images Using Deep Convolutional Neural Networks*. [online] arXiv.org. Available at: https://arxiv.org/abs/1710.00886 [Accessed 20 Feb. 2019].