Developing a Latent Variable Model for Critical Care Patients in the MIMIC-III Dataset using Variational Autoencoders

Scott Fleming Matt Millett Jesper Westell

scottyf@stanford.edu millett@stanford.edu jesperw@stanford.edu

Motivation and Overview

Question: *Are there distinguishable latent features among ICU patients? Can these be interpreted as key indicators for doctors to use?*

- Electronic Medical Records (EMR) have recently been implemented across hospitals in the United States.
- This lets clinicians and researchers more easily access large amounts of data, and allows for finding patterns in patients that have not been found with traditional analysis methods.
- This research uses the Variational Autoencoder (VAE) to characterize latent features in health data of Intensive Care Unit (ICU) patients.
- Experiments conducted using
  - VAE with a Multivariate Gaussian Prior
  - VAE with a Multivariate Gaussian Mixture Prior (GMVAE)

MIMIC-III Database & Dataset

- De-identified EMR data for about ~43,000 adult patients
- Patients admitted to the Beth Israel Deaconess Medical Center ICU between 2001-2012
- Record-level data includes demographics, hourly vital measurements, lab test results, medications, mortality, ICD9 diagnoses, and discharge summaries.

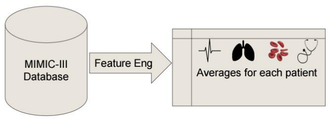

Data Preprocessing

- Data is ugly- lots of NaN's, some patients are 178 years old.
- There's a record for every nurse or physician entry for every patient, across multiple ICU stays per patient.
- Used last-value-carried-forward imputation to fill in missing chart information; used KNN imputation to fill in the rest.
- Grouped chart events by ICU stay, took the average of each vital to get a crude snapshot of the patient's visit.
- Did the same thing for an outcome of potential interest, the SIRS criterion.

Traditional VAE

- Originally proposed by Kingma et al. (2013) [1].
- Encodes input data, **x**, into a latent representation, **z**, adds noise to the latent representation, and decodes **z** back into the original data space.
- **z** is pushed to form a Multivariate Gaussian distribution with a diagonal covariance matrix.
- Utilizes the generative model

$$
\begin{aligned}
p_\theta(x, z) &= p_\theta(z)p_\theta(x \mid z) \\
z &\sim \mathcal{N}(\mu_z, \sigma_z^2) \\
x &\sim \mathcal{N}(\mu_x(z), \sigma_x^2(z))
\end{aligned}
$$

and inference model,

$$
\begin{aligned}
&q_\phi(z \mid x) \\
&z \sim \mathcal{N}(\mu_z(x), \sigma_z^2(x))
\end{aligned}
$$

GMVAE

- Proposed by Shu (2016) [2].
- To explore the possibility that there may be underlying distinctions between groups of patients, we have utilized a Gaussian Mixture Prior for the VAE.
- Hence, **z** is now pushed to form a Multivariate Gaussian Mixture distribution with diagonal covariance matrices.
- Utilizes the generative model

$$
\begin{aligned}
p_\theta(x, y, z) &= p_\theta(y)p_\theta(z \mid y)p_\theta(x \mid z) \\
y &\sim Cat\left(\frac{1}{K}\right) \\
z &\sim \mathcal{N}(\mu_z(y), \sigma_z^2(y)) \\
x &\sim \mathcal{N}(\mu_x(z), \sigma_x^2(z))
\end{aligned}
$$

and inference model,

$$
\begin{aligned}
q_\phi(y, z \mid x) &= q_\phi(y \mid x)q_\phi(z \mid x, y) \\
y &\sim \text{Multinomial}(\theta(x)) \\
z &\sim \mathcal{N}(\mu_z(x, y), \sigma_z^2(x, y))
\end{aligned}
$$

- Model architecture:

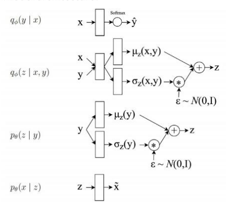

Choosing Hyperparameters

- Used the Calinski-Harabasz Index in determining the optimal number of clusters.
- Ratio of the between-clusters dispersion mean and within cluster dispersion.

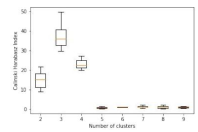

- Grid search on number of latent variables and number of hidden layer variables

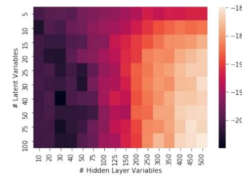

- Tried He initialization.
- Included batch normalization layer into the architecture.
- Implemented early stopping with min_delta=0.0001 and patience=10.
- KL Divergence term is a natural regularizer.
- Using Pairwise Stability Score to measure consistency between multiple runs, and between GMVAE and simpler model such as PCA + GMM clustering.

Future Work

- Find more cluster-friendly features. Hopkins statistic: 0.502.
- Increase GMVAE stability: currently doesn't split into clusters on our dataset, is nondeterministic. (May be due to randomness in MIMIC dataset).
- Need to more deeply explore objective optimization metrics
- Are there other clinical metrics besides just SIRS criteria that we could relate to the latent variable representation?
- How would our model compare to a multitask neural network in predicting outcomes of interest?
- Finding a good way to compute pairwise stability of GMVAE clusters across multiple runs of the algorithm.

Results

- Our GMVAE clustered successfully on MNIST, but not on MIMIC-III.

MNIST sampling

GMVAE clustering t-SNE

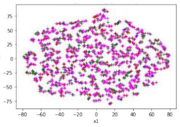

- Latent Variable correlations appear to be clearer in our GMVAE than in a traditional VAE.

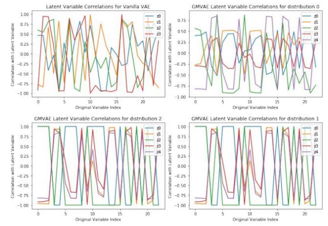

Discussion

- Finding meaningful latent features in the MIMIC-III Dataset using a GMVAE is a hard task.
- Medical data is especially difficult to work with.
- Training the GMVAE on the MNIST dataset shows promise in terms of clustering solutions, but harder to confirm clusterings on MIMIC-III.
- Distinguishable latent features may be easier to find in image data, compared to medical data.

References and Acknowledgements

1. D. P. Kingma and M. Welling. Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114, 2013.
2. R. Shu. Gaussian mixture vae: Lesson in variational inference, generative models, and deep nets. http://ruishu.io/2016/12/25/gmvae/, 2016.