Epitomic variational graph autoencoder

Rayyan Ahmad Khan, Muhammad Umer Anwaar and Martin Kleinsteuber
Technical University of Munich & Mercateo AG, Munich, Germany

Abstract

The learning capacity of variational autoencoders is limited by **over-pruning** - a phenomenon where many latent variables fail to capture useful information about the input data and the corresponding hidden units become **inactive**. This adversely affects learning diverse and interpretable latent representations. As variational graph autoencoder (VGAE) extends VAE for graph-structured data, it inherits the over-pruning problem. In this paper, we adopt a model based approach and propose epitomic VGAE (EVGAE), a generative variational framework for graph datasets which successfully mitigates the over-pruning problem and also boosts the generative ability of VGAE. We consider EVGAE to consist of multiple sparse VGAE models, called epitomes, that are groups of latent variables sharing the latent space. This approach aids in increasing active units as epitomes compete to learn better representation of the graph data. We verify our claims via experiments on three benchmark datasets. Our experiments show that EVGAE has a better generative ability than VGAE. Moreover, EVGAE outperforms VGAE on link prediction task in citation networks.

Unit Activity

Intuition:
An active unit has different values for different inputs.

Definition: $$A_u = \mathrm{Cov}_x(\mathbb{E}_{u\sim q(u|x)}[u])$$

A unit u is said to be active if $A_u > 0.02$

VGAE Vs Pure VGAE vs EVGAE

Loss function of *Pure* VGAE is given by:

$$L(A, \overline{A}) + D_{KL}(\mathcal{N}(\mu,\sigma)||\mathcal{N}(0,1))$$

where **A** is the adjacency matrix. The first term is the reconstruction loss, given by BCE(binary cross-entropy) between input and reconstructed edges. The second term is the KL-divergence between the learnt Gaussian distribution parameters and the standard Gaussian.

Loss function of VGAE introduces a weight to KL-term:

$$L(A, \overline{A}) + \beta D_{KL}(\mathcal{N}(\mu,\sigma)||\mathcal{N}(0,1)); \quad \beta = \frac{1}{N}$$

This results in better reconstruction at the cost of poor generativeness.

Loss function of EVGAE is given by:

$$L = \overbrace{\mathrm{BCE}}^{L_1} + \overbrace{\sum_{i=1}^{N} D_{KL}\Big(\mathrm{Cat}(\pi_i(\mathcal{G}))||\,\mathcal{U}(1,M)\Big)}^{L_2}$$
$$+ \underbrace{\sum_{i=1}^{N}\sum_{y_i}\pi_i(\mathcal{G})\sum_{j=1}^{D}E[y_i,j]D_{KL}\Big(\mathcal{N}\big(\mu_i^j(\mathcal{G}),(\sigma_i^2)^j(\mathcal{G})\big)||\mathcal{N}(0,1)\Big)}_{L_3}.$$

EVGAE Architecture

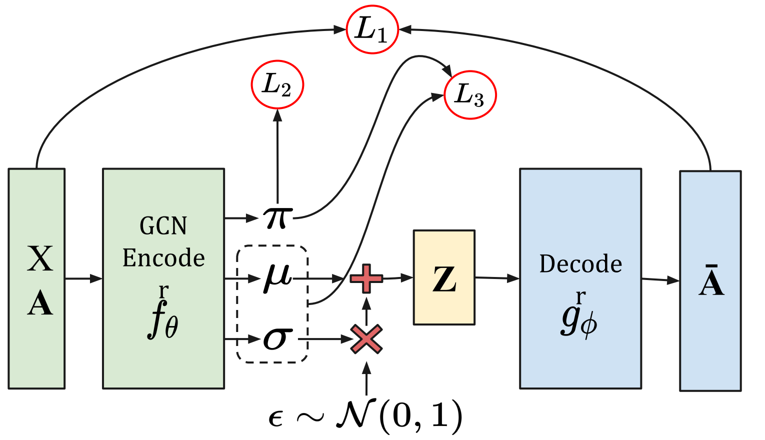

Results

KL-divergence - Cora **Unit Activity - Cora**

Pure VGAE

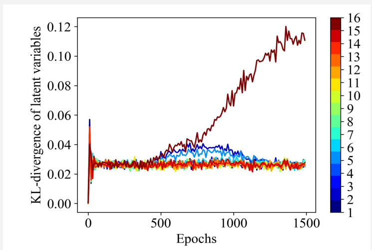

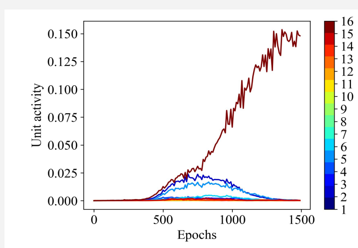

VGAE

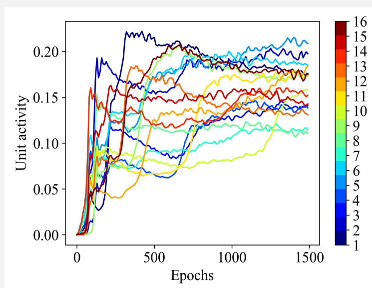

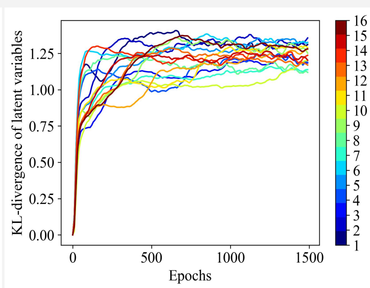

EVGAE

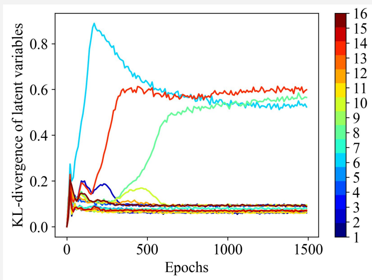

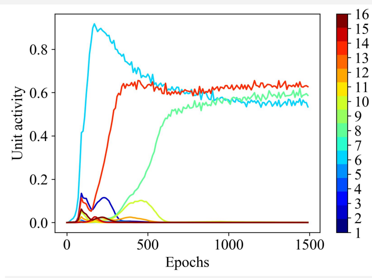

EVGAE as a compromise between VGAE and Pure VGAE