Bayesian Semisupervised Learning with Deep Generative Models

Jonathan Gordon, José Miguel Hernández-Lobato
Department of Engineering University of Cambridge

Objectives

The project is concerned with developing a deep generative model (DGM) for semi-supervised and active learning. The main objectives are:

- Develop a DGM with a discriminative component
- Accommodate semi-supervised learning
- Extend to Bayesian training and active learning

Variational Inference in Deep Models

Variational inference has enabled efficient training of large scale generative models. An example is the Variational autoencoder (VAE) (depicted for semi-supervised learning).

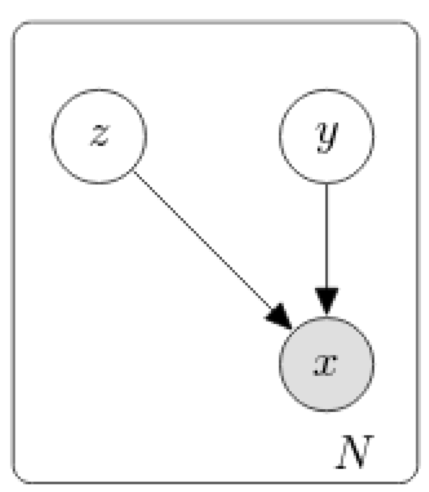

A parameterized approximate posterior $q_\phi(z|x)$ is introduced, and the ELBO is maximized. In the case of a VAE, we have:

$$\mathcal{L}_{vae} = \mathbb{E}_{q_\phi}\big[\log p_\theta(x|z)\big] - \mathcal{D}_{kl}\Big(q_\phi(z|x)\|p_\theta(z)\Big) \approx \frac{1}{L}\sum_{l=1}^{L} \log p_\theta(x|z^{(l)}) - \log q_\phi(z^{(l)}|x) + \log p_\theta(z^{(l)}) \tag{1}$$

The estimator can be approximated using the reparameterization trick:

$$z^{(l)} = g_\phi(\epsilon^l, x) \sim q_\phi(z|x), \quad \epsilon \sim p(\epsilon) \tag{2}$$

Despite excellent performance in semi-supervised tasks, use of VAEs adopt inference networks as classifiers, discarding the model after training.

Contact Details

J. Gordon: jg801@cam.ac.uk
J. M. Hernández-Lobato: jmh223@cam.ac.uk

Model Description

We propose the following model:

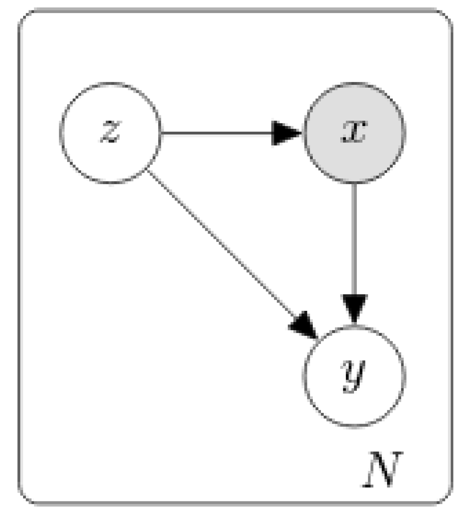

With a VAE $p_\theta(z|x)$ for unsupervised learning and a Bayesian neural network $p_\theta(y|x,z)$ for discriminative tasks.

$$\begin{aligned} p(z) &= \mathcal{N}(z; \mathbf{0}, \boldsymbol{I}) \\ p_\theta(x|z) &= \mathcal{N}(x; \mu(z), \nu(z)) \\ p_\theta(y|z,x) &= \mathrm{Cat}(y; \pi(x,z)) \end{aligned} \tag{3}$$

where the distributions of x, y are parameterized with deep neural networks.

Variational Training

We introduce the approximate posterior network:

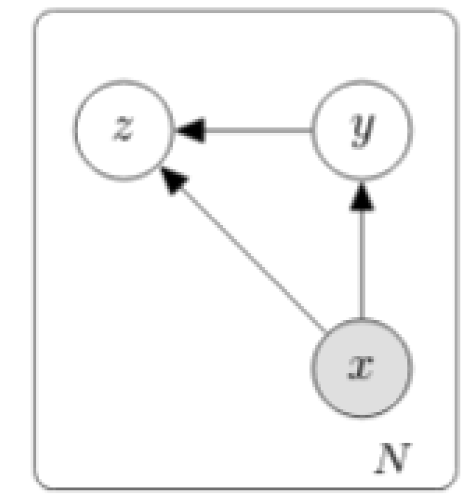

$$q_\phi(z,y|x) = q_\phi(z|x,y)q_\phi(y|x) \tag{4}$$

Yielding an overall ELBO:

$$\mathcal{L}(\theta,\phi;x,y) = \underbrace{\mathcal{L}^l(\theta,\phi;x,y)}_{\text{Labeled}} + \underbrace{\mathcal{L}^u(\theta,\phi;x,y)}_{\text{Unlabeled}} \tag{5}$$

Expectations w.r.t. $q_\phi(y|x)$ can be taken with exact enumeration. Predictions are made with Gibbs sampling.

Initial Experimental Results

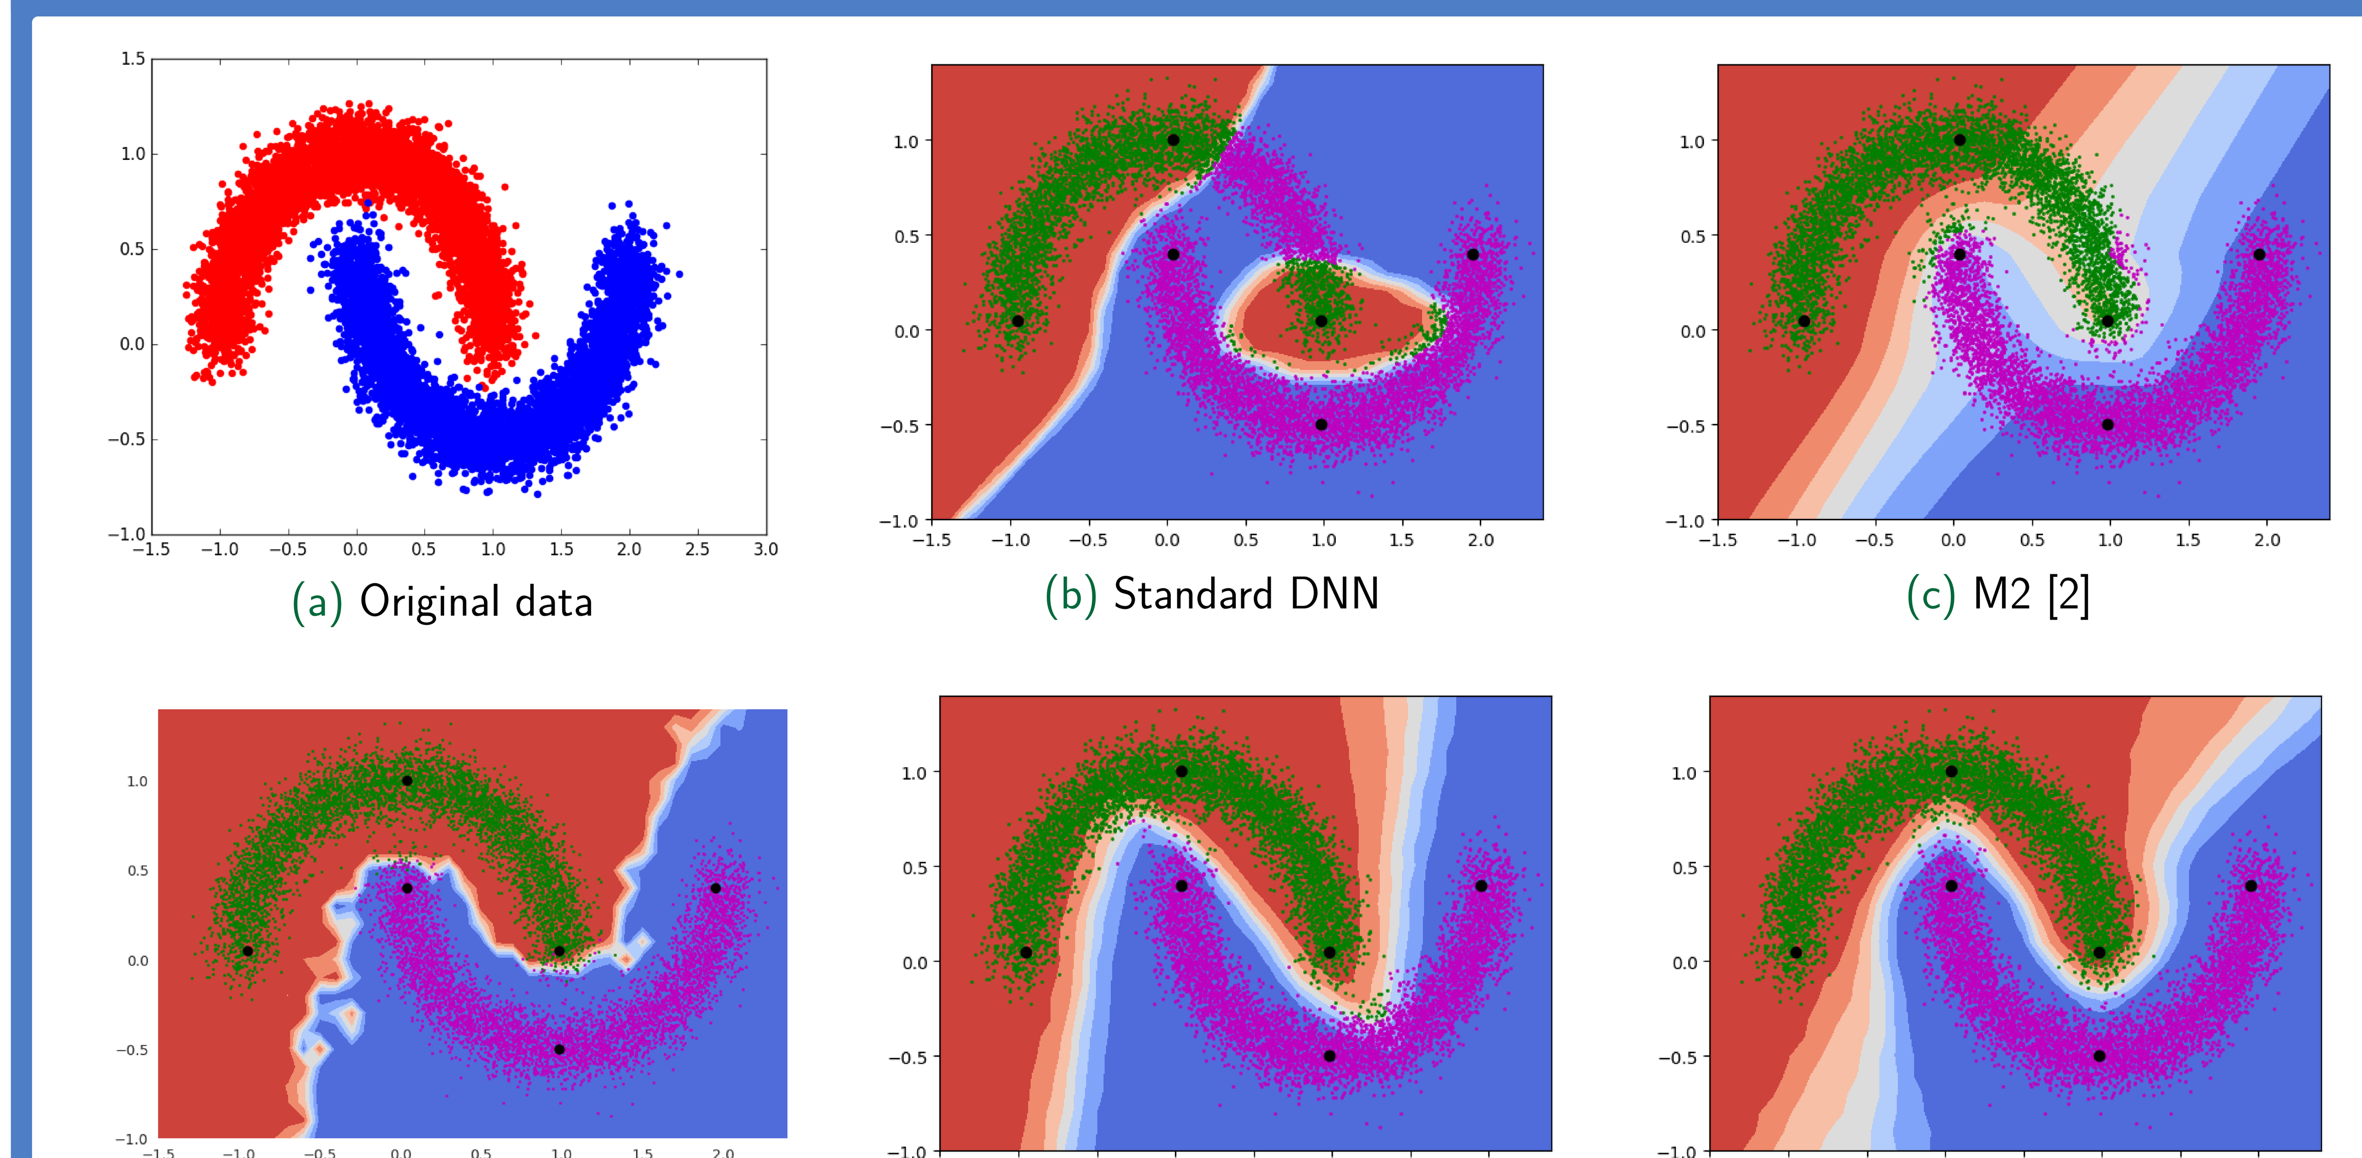

(a) Original data
(b) Standard DNN
(c) M2 [2]
(d) SDGM [4]
(e) SSLPE
(f) SSLAPD

Bayesian Training

We can extend the model to Bayesian training. The ELBO becomes:

$$\begin{aligned} \mathcal{L}^l(\theta,\phi;x,y) =& \mathbb{E}_{q_\phi}[\log p_\theta(x,y|z,w)] \\ & - \mathcal{D}_{KL}(q_\phi(z|x,y)\|p(z)) \\ & - \mathcal{D}_{KL}(q_\phi(w)\|p(w)) \end{aligned} \tag{6}$$

where we have introduced the factorized approximation:

$$q(w) = \prod \mathcal{N}(w_j; \mu_{w_j}, \sigma^2_{w_j}) \tag{7}$$

We can use the reparameterization trick for sampling $w^l \sim q(w)$, and optimize as previously. Accounting for model uncertainty opens the door to active learning.

Future Work

The model has been implemented and experimented with in limited settings. Future goals are:

- Improve performance in the semi-supervised setting
- Stabilize Bayesian training
- Experiment with benchmark datasets
- Implement and explore active learning schemes

References

[1] Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[2] Diederik P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. In *Advances in Neural Information Processing Systems*, pages 3581–3589, 2014.

[3] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks. *arXiv preprint arXiv:1505.05424*, 2015.

[4] Lars Maaløe, Casper Kaae Sønderby, Søren Kaae Sønderby, and Ole Winther. Auxiliary deep generative models. *arXiv preprint arXiv:1602.05473*, 2016.