

# UNSUPERVISED LEARNING WITH IMBALANCED DATA VIA STRUCTURE CONSOLIDATION LATENT VARIABLE MODEL

**Fariba Yousefi<sup>1</sup>, Zhenwen Dai<sup>2</sup>, Carl Henrik Ek<sup>3</sup>, Neil Lawrence<sup>4</sup>**

<sup>1,4</sup> University of Sheffield, Sheffield, UK

<sup>2</sup> Inferentia Ltd, UK

<sup>3</sup> University of Bristol, Bristol, UK

<sup>1,4</sup> {f.yousefi, n.lawrence}@sheffield.ac.uk

<sup>2</sup> zhenwendai@gmail.com

<sup>3</sup> carlhenrik.ek@bristol.ac.uk

## ABSTRACT

Unsupervised learning on imbalanced data is challenging because, when given imbalanced data, current model is often dominated by the major category and ignores the categories with small amount of data. We develop a latent variable model that can cope with imbalanced data by dividing the latent space into a shared space and a private space. Based on Gaussian Process Latent Variable Models, we propose a new kernel formulation that enables the separation of latent space and derive an efficient variational inference method. The performance of our model is demonstrated with an imbalanced medical image dataset.

## 1 INTRODUCTION

In many medical applications, e.g. pathology, negatively labelled data is extremely easy to obtain (e.g. healthy cells). Positive labels, on the other hand, can be harder to acquire (e.g. particular disease morphologies). These massively unbalanced problems are challenging for most algorithms because the negative class tends to dominate the objective function and the resulting model performs poorly. In practice it is often better to throw away much of the negative data and rebalance the data set.

Unsupervised learning has been attracting a lot of attention as it has the potential to serve as an underpinning technology for a range of challenges such as generative modeling, missing data imputation and coping with multiple data modalities. Unsupervised learning can also be applied to a wider range of data sets, because it does not rely on having carefully labelled data available.

In this paper we explore the possibility of using a variant unsupervised learning algorithm to solve the problem of label balance. We build latent variable models that can simultaneously accommodate a very large number of negative examples, sharing their characteristics appropriately with the positive class, while simultaneously allowing the model to characterise the manner in which the positive class is differently characterised through preserved (or private) latent spaces that are separately learned for each class. The resulting model does not suffer from the standard challenges in this domain. We compare with a variant of the discriminative GP-LVM (the model that underpinned GaussianFace) and show significantly improved performance.

Our probabilistic latent variable model divides its latent space into a shared space of all the categories and a private space for each category (Damianou et al., 2012). The shared space accounts for capturing the common regularities among categories (e.g. positive and negative class) and the private space is dedicated to model the variance specific to individual categories. Because the modelling of the private space is category specific, there is no domination of its characteristics by the larger category. Thus the data in each category can be modeled appropriately while the common regularities are still exploited.

We implement the idea of shared and private space in the framework of Gaussian Process latent variable models (GPLVM, Lawrence, 2005) by deriving a particular covariance function (kernel) that enables such separation. We exploit closed form variational lower bounds of the log marginal likelihood of the proposed model, which to provide an efficient approximation inference method.

The performance of our model is evaluated with a real image dataset, in which the positive and negative data are extremely imbalanced. We show our model still can learn from imbalanced data and perform well in both generative and discriminative tasks.

## 2 STRUCTURE CONSOLIDATION LATENT VARIABLE MODEL

We assume the dataset is represented as a set of fixed length vectors  $\mathbf{Y} \in \mathbb{R}^{N \times D}$ , where  $N$  is the number of data points and  $D$  is the dimensionality of individual data points. Additionally, a label of category is associated with each data point,  $\mathbf{c} = (c_1, \dots, c_N)$ ,  $c_i \in \{1, \dots, C\}$ , where  $C$  indicates the number of categories in the dataset. We aim at building a probabilistic model  $p(\mathbf{Y})$  that is robust when the numbers of data in different categories are highly imbalanced.

We assume the data associated with a set of latent representations  $\mathbf{X} \in \mathbb{R}^{N \times Q}$ , where  $Q$  is the dimensionality of the latent space. The latent representations are related to the observed data through an unknown mapping function  $f$  and  $f$  follows a prior distribution that is defined as a Gaussian process,

$$\mathbf{y} = f(\mathbf{x}) + \epsilon, \quad f \sim \mathcal{GP}(0, k), \quad (1)$$

where  $\epsilon \in \mathcal{N}(0, \sigma^2 \mathbf{I})$  denotes the observation noise and  $k$  is the kernel function. Given the observed data  $\mathbf{Y}$ , we wish to obtain a posterior estimate for both the latent representation  $\mathbf{X}$  and the unknown mapping function  $f(\cdot)$ . In our model we separate the latent space into a shared space with the dimensionality  $Q_s$  and a private space with the dimensionality  $Q_p$ . Therefore, a latent representation can be denoted as  $\mathbf{x} = [\mathbf{x}_s^\top, \mathbf{x}_p^\top]^\top$ ,  $\mathbf{x}_s \in \mathbb{R}^{Q_s}$ ,  $\mathbf{x}_p \in \mathbb{R}^{Q_p}$ , where  $\mathbf{x}_s$  and  $\mathbf{x}_p$  are the latent representations in shared and private space respectively. With the separated latent representation, we define the kernel function in our model as

$$k((\mathbf{x}, c_x), (\mathbf{x}', c_{x'})) = k_s(\mathbf{x}_s, \mathbf{x}'_s) + k_p((\mathbf{x}_p, c_x), (\mathbf{x}'_p, c_{x'})), \quad (2)$$

where  $k_s$  is the kernel function for the shared space and  $k_p$  is the kernel function for the private space. The shared kernel can be any kernel function built on a vector space from the literature. However, the private kernel is defined to take the following form:

$$k_p((\mathbf{x}_p, c_x), (\mathbf{x}'_p, c_{x'})) = \begin{cases} k'((\mathbf{x}_p, c_x), (\mathbf{x}'_p, c_{x'})), & c_x = c_{x'}, \\ 0, & c_x \neq c_{x'}, \end{cases} \quad (3)$$

where  $k'$  is the kernel function chosen to calculate the covariance and  $c_x$  is the label of category for the data point  $\mathbf{x}$ . We give a unit Gaussian prior distribution to latent representations  $\mathbf{x} \sim \mathcal{N}(0, \mathbf{I})$ . The log marginal likelihood for the proposed model can be derived as  $\log p(\mathbf{Y}) = \log \int p(\mathbf{Y}|\mathbf{X})p(\mathbf{X})d\mathbf{X}$ . There is no analytical solution for this marginal likelihood. We apply variational inference and derive a closed form lower bound of the log marginal likelihood, by following a sparse Gaussian process approximation (Titsias & Lawrence, 2010):

$$\log p(\mathbf{Y}) \geq \sum_{d=1}^D \tilde{F}_d(q) - \text{KL}(q(\mathbf{X})||p(\mathbf{X})), \quad (4)$$

$$\tilde{F}_d(q) = \log \left[ \frac{|\mathbf{K}_{uu}|^{\frac{1}{2}}}{(2\pi\sigma^2)^{\frac{N}{2}} |\beta\Psi_2 + \mathbf{K}_{uu}|^{\frac{1}{2}}} e^{-\frac{1}{2}\mathbf{y}_d^\top W \mathbf{y}_d} \right] - \frac{\psi_0}{2\sigma^2} + \frac{1}{2\sigma^2} \text{Tr}(\mathbf{K}_{uu}^{-1}\Psi_2), \quad (5)$$

where  $W = \frac{1}{\sigma^2} \mathbf{I} - \frac{1}{\sigma^4} \Psi_1 (\frac{1}{\sigma^2} \Psi_2 + \mathbf{K}_{uu})^{-1} \Psi_1^\top$ , and  $\psi_0, \Psi_1, \Psi_2$  are the expectation of covariance matrices w.r.t. the variational posterior  $q(\mathbf{X})$ . In our model, these expectation are derived as

$$\psi_0 = \sum_{n=1}^N \left\langle k_s(\mathbf{x}_s^{(n)}, \mathbf{x}_s^{(n)}) \right\rangle_{q(\mathbf{x}_s^{(n)})} + \left\langle k'(\mathbf{x}_p^{(n)}, \mathbf{x}_p^{(n)}) \right\rangle_{q(\mathbf{x}_p^{(n)})} \quad (6)$$

$$(\Psi_1)_{nm} = \left\langle k_s(\mathbf{x}_s^{(n)}, \mathbf{z}_s^{(m)}) \right\rangle_{q(\mathbf{x}_s^{(n)})} + \left\langle k_p(\mathbf{x}_p^{(n)}, \mathbf{z}_p^{(m)}) \right\rangle_{q(\mathbf{x}_p^{(n)})} \quad (7)$$

$$\begin{aligned} (\Psi_2)_{mm'} &= \sum_{n=1}^N \left\langle k_s(\mathbf{x}_s^{(n)}, \mathbf{z}_s^{(m)}) k_s(\mathbf{x}_s^{(n)}, \mathbf{z}_s^{(m')}) \right\rangle_{q(\mathbf{x}_s^{(n)})} + \left\langle k_p(\mathbf{x}_p^{(n)}, \mathbf{z}_p^{(m)}) k_p(\mathbf{x}_p^{(n)}, \mathbf{z}_p^{(m')}) \right\rangle_{q(\mathbf{x}_p^{(n)})} \\ &\quad + \left\langle k_s(\mathbf{x}_s^{(n)}, \mathbf{z}_s^{(m)}) \right\rangle_{q(\mathbf{x}_s^{(n)})} \left\langle k_p(\mathbf{x}_p^{(n)}, \mathbf{z}_p^{(m')}) \right\rangle_{q(\mathbf{x}_p^{(n)})} \\ &\quad + \left\langle k_p(\mathbf{x}_p^{(n)}, \mathbf{z}_p^{(m)}) \right\rangle_{q(\mathbf{x}_p^{(n)})} \left\langle k_s(\mathbf{x}_s^{(n)}, \mathbf{z}_s^{(m')}) \right\rangle_{q(\mathbf{x}_s^{(n)})} \end{aligned} \quad (8)$$

where  $\mathbf{z}$  is the variational parameter known as inducing inputs.

### 3 EXPERIMENT

Mitosis detection is a stage in tumour assessment that involves determining whether individual cells are in mitosis (dividing to reproduce). These cells are rare. We used data from the assessment of mitosis detection algorithm 2013 (AMIDA13, Veta et al., 2015) challenge which is publicly available. The main goal of the challenge is to find proper mitosis detection methods that can be automatic or semi-automatic. We use the training set from the challenge, which consists of tissue images of 12 patients and are annotated by human experts. We preprocess the tissue images with the algorithm by Snell (2013) and focus on the generated candidate image patches. The resulting image set contains 146,562 grey-scale image patches ( $30 \times 30$  pixels), of which 550 are positive (mitosis) according to manual annotation. We randomly take 80% of positive images and 5,000 negative images as the training data. This gives in total 5,440 images. Some examples of the training data is shown in Figure 2a. We apply for SCLVM to this dataset and used an exponentiated quadratic kernel for both the shared and private space and set the dimensionality of both shared and private space to five.

Both the latent representations and kernel parameters are optimized until convergence. The resulting latent space is visualized in Figure 1. The positive and negative images present similar structures in the shared space, which demonstrates the discovered common regularities, and their private space are significantly different from each other. To demonstrate the ability of SCLVM in balancing the modeling capabilities between imbalanced categories, we draw samples from the learned latent space of SCLVM for both positive and negative categories (see Figure 2b). The generated samples from positive and negative categories are clearly different from each other and they capture some characteristics of their own categories. We further evaluate the learned latent space by performing classification on test set (the rest of positive examples plus randomly sampled negative examples, in total 1000 images). We compare SCLVM with BGPLVM (Titsias & Lawrence, 2010) and DG-PLVM<sup>1</sup> (Urtasun & Darrell, 2007) with ten test sets. We apply weighted SVM with an exponentiated quadratic kernel on the latent space from the BGPLVM and DGPLVM. The results are shown in Table 1. Note that BGPLVM requires an additional classification model to be learned. This does not provide probabilities over the classes other than in the ad-hoc manner that an SVM will. Similarly DGPLVM does learn a space that reflects the class information, but it does not provide means to get the posterior over the classes. Our model is the only one that learns the classification jointly with the model and provides a principled way of getting probabilities over the classes.

### 4 CONCLUSION

We presented a probabilistic latent variable model that can cope with imbalanced data. We developed a kernel that separates the latent space into a shared space and a private space. An efficient variational inference method is proposed by deriving a closed form lower bound of marginal likelihood. Beyond

<sup>1</sup>Due to its complexity  $O(N^3)$ , we only use 1000 images for training (440 positive and 560 negative).




Figure 1: The visualization of the training data in learned latent spaces. The first figure shows the positive and negative data in two of the shared dimensions. The second and third figures show the two of the private dimensions for the negative and positive data respectively. The fourth figure shows the learned latent space from BGPLVM.




Figure 2: (a) Some examples in the data sets. (b) Samples generated from the trained SCLVM. In both figures, the first two rows corresponds to positive images and the last two rows corresponds to negative images.

Table 1: Classification performance. The mean and standard deviation from ten test sets are shown.

	SCLVM	BGPLVM (SVM)	DGPLVM (SVM)
precision	$0.426 \pm 0.024$	$0.306 \pm 0.013$	$0.242 \pm 0.008$
recall	$0.555 \pm 0.007$	$0.827 \pm 0.000$	$1.000 \pm 0.000$
F1 score	$0.482 \pm 0.015$	$0.447 \pm 0.014$	$0.390 \pm 0.011$

the shown example, the ability of jointly modeling multiple data categories and handling imbalanced datasets can be linked to many other areas such as transfer learning.

## REFERENCES

- Andreas Damianou, Carl Henrik Ek, Michalis K. Titsias, and Neil D. Lawrence. Manifold relevance determination. In John Langford and Joelle Pineau (eds.), *Proceedings of the International Conference in Machine Learning*, volume 29, San Francisco, CA, 2012. Morgan Kaufman.
- Neil D. Lawrence. Probabilistic non-linear principal component analysis with Gaussian process latent variable models. *Journal of Machine Learning Research*, 6:1783–1816, 11 2005.
- Violet Snell. *Shape and Texture Recognition for Automated Analysis of Pathology Images*. PhD thesis, Centre for Vision, Speech and Signal Processing, University of Surrey, Surrey, UK, 2013.
- Michalis K. Titsias and Neil D. Lawrence. Bayesian Gaussian process latent variable model. In Yee Whye Teh and D. Michael Titterton (eds.), *Proceedings of the Thirteenth International Workshop on Artificial Intelligence and Statistics*, volume 9, pp. 844–851, Chia Laguna Resort, Sardinia, Italy, 13-16 May 2010. JMLR W&CP 9.
- Raquel Urtasun and Trevor Darrell. Discriminative gaussian process latent variable model for classification. In *Proceedings of the 24th international conference on Machine learning*, pp. 927–934. ACM, 2007.

Mitko Veta, Paul J Van Diest, Stefan M Willems, Haibo Wang, Anant Madabhushi, Angel Cruz-Roa, Fabio Gonzalez, Anders BL Larsen, Jacob S Vestergaard, Anders B Dahl, et al. Assessment of algorithms for mitosis detection in breast cancer histopathology images. *Medical image analysis*, 20(1):237–248, 2015.