
General Latent Feature Modeling for Data Exploration Tasks

Isabel Valera¹ Melanie Fernandez Pradier² Zoubin Ghahramani^{1,3}

Abstract

This paper introduces a general Bayesian non-parametric latent feature model suitable to perform automatic exploratory analysis of heterogeneous datasets, where the attributes describing each object can be either discrete, continuous or mixed variables. The proposed model presents several important properties. First, it accounts for heterogeneous data while can be inferred in linear time with respect to the number of objects and attributes. Second, its Bayesian non-parametric nature allows us to automatically infer the model complexity from the data, i.e., the number of features necessary to capture the latent structure in the data. Third, the latent features in the model are binary-valued variables, easing the interpretability of the obtained latent features in data exploration tasks.

1. Introduction

Latent feature models allow us to compact in a few features the immense redundant information present in the observed data, by capturing the statistical dependencies among the different objects and attributes. As a consequence, they appear as suitable tools to perform data exploratory analysis, i.e, they may help us to better understand the data (Blanco et al., 2013; Ruiz et al., 2013).

There is an extensive literature in latent feature modeling of homogeneous data, where all the attributes describing each object in the database present the same (continuous or discrete) nature. In particular, these works assume that databases contain only either continuous data, usually modeled as Gaussian variables (Griffiths & Ghahramani, 2011; Todeschini et al., 2013), or discrete, that can be either modeled by discrete likelihoods (Li, 2009; Ruiz et al., 2013; Gopalan et al., 2014) or simply treated as Gaussian variables (Blanco et al., 2013; Todeschini et al., 2013). However, there still exists a lack of works dealing with heteroge-

neous databases, which in fact are common in real applications. As motivating examples, Electronic Health Records from hospitals might contain lab measurements (often positive real-valued or real-valued data), diagnoses (categorical data) and genomic information (ordinal, count data and categorical data); also, a survey often contains diverse information about the participants such as age (count data), gender (categorical data), salary (positive real data), etc. Despite this diversity of data types, the standard approach when dealing with heterogeneous datasets is to treat all the attributes, either continuous or discrete, as Gaussian variables.

This paper presents a general latent feature model (GLFM) suitable for heterogeneous datasets, where the attributes describing each object can be either discrete, continuous or mixed variables. Specifically, we account for real-valued and positive real-valued as examples of continuous variables, and categorical, ordinal and count data as examples of discrete variables. The proposed model extends the essential building block of Bayesian nonparametric latent feature models, the Indian Buffet Process (IBP) by (Griffiths & Ghahramani, 2011), to account for heterogeneous data while maintaining the model complexity of conjugate models. Among all the available latent feature models in the literature, we opt for the IBP due to two main reasons. First, the nonparametric nature of the IBP allows us to automatically infer the appropriate model complexity, i.e., the number of necessary features, from the data. Second, the IBP considers binary-valued latent features which has been shown to provide more interpretable results in data exploration than standard real-valued latent feature models (Ruiz et al., 2012; 2013). The standard IBP assumes real-valued observations combined with conjugate likelihood models, allowing for fast inference algorithms (Doshi-Velez & Ghahramani, 2009). However, we here aim at dealing with heterogeneous databases, such that conjugacy might not be straightforwardly available.

In order to propose a general observation model for the IBP that accounts for heterogeneous data while keeping the properties of conjugate models, we exploit two key ideas. First, we introduce an auxiliary real-valued variable (also called *pseudo-observation*), such that, conditioned on it, the model behaves as the standard linear-Gaussian IBP in (Griffiths & Ghahramani, 2011). Second, we assume that there exists a function that transforms the pseudo-

¹University of Cambridge, Cambridge, United Kingdom;

²Universidad Carlos III de Madrid, Spain ³Uber AI Labs, San Francisco, California, USA. Correspondence to: Isabel Valera <miv24@cam.ac.uk>.




Figure 1. Illustration of the GLFM.

observation into the actual observation, mapping the real line into the (discrete or continuous) observation space of each attribute in the data. These two key ideas allow us to derive an efficient inference algorithm based on collapsed Gibbs sampling, which presents linear complexity with the number of objects and attributes in the data.

Our experiments provide examples of how to use the proposed model for data exploration in real-world datasets. Additionally, a software library implementing the GLFM, as well as the necessary scripts to perform automatic data exploration, is publicly available at <https://github.com/ivaleraM/GLFM>.

2. General Latent Feature Model

We introduce the GLFM, which is a general Bayesian non-parametric latent feature model suitable for data exploration of heterogeneous datasets, where the attributes describing each object can be either discrete, continuous or mixed variables. Specifically, the GLFM accounts for the following data types:

- Continuous variables:
 1. Real-valued, i.e., $x_n^d \in \mathfrak{R}$
 2. Positive real-valued, i.e., $x_n^d \in \mathfrak{R}_+$.
- Discrete variables:
 1. Categorical data, i.e., x_n^d takes a value in a finite unordered set, e.g., $x_n^d \in \{\text{‘blue’}, \text{‘red’}, \text{‘black’}\}$.
 2. Ordinal data, i.e., x_n^d takes values in a finite ordered set, e.g., $x_n^d \in \{\text{‘never’}, \text{‘sometimes’}, \text{‘often’}, \text{‘usually’}, \text{‘always’}\}$.
 3. Count data, i.e., $x_n^d \in \{0, \dots, \infty\}$.

The GLFM builds on the IBP (Griffiths & Ghahramani, 2011), and therefore, it assumes that each observation x_n^d can be explained by a potentially infinite-length binary vector \mathbf{z}_n whose elements indicate whether a latent feature is active or not for the n -th object; and a (real-valued) weighting vector \mathbf{B}^d , whose elements weight the influence of each latent feature in the d -th attribute¹. Since the product of the latent feature vector and the weighting vector leads to a real-valued variable, it is necessary to map this variable to the desirable output (continuous or discrete) space, for example, the positive real line or the

finite ordered set $\{\text{‘low’}, \text{‘medium’}, \text{‘high’}\}$. Thus, the GLFM assumes the existence of intermediate real-valued auxiliary variables $y_n^d \sim \mathcal{N}(\mathbf{z}_n^T \mathbf{B}^d, \sigma_d^2)$, called *pseudo-observation*, and a transformation function $f_d(\cdot)$ that maps this variable into the actual observation x_n^d , i.e., $x_n^d = f_d(y_n^d + u)$ where $u \sim \mathcal{N}(0, \sigma_u^2)$ is an auxiliary noise with zero mean and small variance σ_u^2 . Additionally, the GLFM accounts for a bias term similar to the one in (Ruiz et al., 2012; 2013), which corresponds to an extra latent feature that is active for every object in the data and eases the interpretability of the latent features, as shown in next section.

Figure 1 illustrates the GLFM by showing the corresponding graphical model together with an example of the generative model for an ordinal attribute taking values in the ordered set $\{\text{‘low’}, \text{‘medium’}, \text{‘high’}\}$. The inference of the GLFM is performed using collapsed Gibbs sampling, which presents linear complexity with respect to the number of objects N and the number of attributes D . Additional details on the model, as well as on the inference algorithm can be found in (Valera et al., 2017).

3. Data Exploration

The main goal of this section is to provide showcase examples about how to include the specific domain knowledge into the proposed GLFM to find and analyze the latent structure underlying data in different application domains, i.e., to perform data exploratory analysis. In particular, we here show examples of how to select the input data for the GLFM, as well as how to enter these data into the model, in order to obtain interpretable results that can be used to get a better understanding of the data.

3.1. Drug effect in a clinical trial for prostate cancer

Clinical trials are conducted to collect data regarding the safety and efficacy of a new drug before it can be sold in the consumer market, if ever. Concretely, the main goal of clinical trials is to prove the efficacy of a new treatment for a disease while ensuring its safety, i.e., check whether its adverse effects remain low enough for any dosage level of the drug. As an example, the publicly available *Prostate Cancer dataset*² collects data of a clinical trial that aimed at analyzing the effects of the drug diethylstilbestrol (DES)

¹For convenience, we here capitalize the vector \mathbf{B}^d .

²<http://biostat.mc.vanderbilt.edu/wiki/Main/DataSets>

General Latent Feature Model

Attribute description	Type of variable
Stage of the cancer	Categorical with 2 categories
DES treatment level	Ordinal with 3 categories
Tumor size in cm ²	Count data
Serum Prostatic Acid Phosphatase (PAP)	Positive real-valued
Prognosis Status (outcome of the disease)	Categorical with 4 categories

Table 1. List of considered attributes for the Prostate Cancer dataset.

as a treatment against prostate cancer. More in detail, the dataset contains information about 502 patients with prostate cancer in stages³ 3 and 4, who entered a clinical trial during 1967-1969 and were randomly allocated to different levels of treatment with DES. The prostate cancer dataset has been used by several studies (Byar & Green, 1980; Lunn & McNeil, 1995) to analyze the survival times of the patients in the clinical trial and the causes behind their death. These studies have pointed out that a large dose of the treatment tends to reduce the risk of a cancer death at any time, but also might result in an increased risk of cardiovascular death.

In this section, we apply the proposed GLFM to the Prostate Cancer dataset to show that the proposed model can be efficiently used to discover the statistical dependencies in the data, which in this example corresponds to the effect of the different levels of treatment with DES in the suffering of prostate cancer and cardiovascular diseases. The prostate cancer dataset consists of 502 patients and 16 attributes, from which we make use of the five attributes listed in Table 1. The selection of these five attributes allows us not only to reduce the number of local minima in the posterior distribution of the proposed model due to the small sample size of the dataset, but also to focus on capturing the statistical dependencies between the target attributes, i.e., the relationship between the different levels of treatment with DES and the suffering of prostate cancer and cardiovascular diseases.

Results. After running our model, we obtain four latent features. Figure 2 shows the effect of the inferred latent features, as well as the bias term, on each dimension/attribute of the data, where we can distinguish two groups of features. The first group accounts for patients in stage 3 and includes the bias term and the 2 first latent features. Within this group, the bias term – or equivalently pattern (0000) – and the first feature – or equivalently pattern (1000) – account for patients in stage 3 with a low average level of treatment with DES (refer to Figure 2b). However, while the bias term models patients with low probability (~ 15%)

³The stage of a cancer describes the size of a cancer and how far it has grown. Stage 3 means that the cancer is already quite large and may have started to spread into surrounding tissues or local lymph nodes. Stage 4 is more severe, and refers to a cancer that has already spread from where it started to another body organ. This is also called secondary or metastatic cancer. Find more details in <http://www.cancerresearchuk.org/about-cancer/what-is-cancer/stages-of-cancer>

of prostate cancer death, the first feature accounts for patients with higher probability (~ 40%) of prostate cancer death, which can be explained by a larger tumor size (refer to Figure 2c). The second feature – or equivalently pattern (0100) – captures patients who exclusively received a high dosage (5 mg) of the drug (refer to Figure 2b). These patients present a small tumor size and the lowest probability of prostatic cancer death, suggesting a positive effect of the drug as treatment for the cancer. However, they also present a significant increase in the probability of dying from a vascular disease (~ 50%), indicating a potential adverse-effect of the drug that increases the risk of suffering from cardiovascular diseases. Such observation is in agreement with previous studies (Byar & Green, 1980; Lunn & McNeil, 1995).

The second group of features corresponds to the activation patterns (0010) and (0001), and accounts for patients in stage 4 with, respectively, mild and severe conditions. In particular, the third feature corresponds to patients with small tumor size, but intermediate values for the PAP biomarker, suggesting a certain spread degree of the tumor compared to the features in the first group, but not as severe as for patients with pattern (0001). Indeed, pattern (0001) models those patients in stage 4 with relatively high tumor size and the highest PAP values—it is thus not surprising that those patients present in turn the highest probability (above 50%) of prostatic death.

3.2. Impact of Social Background on Mental Disorders

Several studies have analyzed the impact of social background in the development of mental disorders (Weissman et al., 1993; Weich & Lewis, 1998). Other studies have focused on finding and analyzing the co-occurring (comorbidity) pattern among the 20 most common psychiatric disorders (Blanco et al., 2013; Ruiz et al., 2013). These studies found that the 20 most common disorders can be divided into three groups: i) externalizing disorders, which include substance use disorders (alcohol abuse and dependence, drug abuse and dependence and nicotine dependence); ii) internalizing disorders, which include mood and anxiety disorders (major depressive disorder (MDD), bipolar disorder and dysthymia, panic disorder, social anxiety disorder (SAD), specific phobia and generalized anxiety disorder (GAD), and pathological gambling (PG)); and iii) personality disorders (avoidant, dependent, obsessive-compulsive (OC), paranoid, schizoid, histrionic and antisocial personality disorders (PDs)). However, up to our knowledge,

there is a lack of work on the impact of social background in the suffering of comorbid disorders.

In this section, we aim at extending the analysis in (Ruiz et al., 2013) to account for the influence of the social background of subjects (such as age, gender, etc.) in the probability of a subject suffering from comorbid disorders. To this end, in addition to the diagnoses of the above 20 psychiatric disorders, we also make use of the information provided by the NESARC, which includes a set of questions on the social background of participants. Specifically, in addition to the diagnoses of the most common 20 psychiatric disorders described above, we include the sex of the participants as input data to the proposed model. We model the gender information of the participants in the NESARC as a categorical variable with two categories: {'male', 'female'}. The percentage of males in the NESARC is approximately 43%. Note also that the diagnoses of the 20 psychiatric disorders correspond to categorical variables with two possible categories, e.g., a patient suffering or not from a disorder.

Results. After running our inference algorithm with the diagnoses of the 20 disorders and the gender of subjects as input data, we obtain three latent features. Figure 3a shows the probability of meeting each diagnostic criteria for the latent feature vectors \mathbf{z}_n listed in the legend and in the database (baseline). Note that the obtained latent features are similar to the ones in (Ruiz et al., 2013), i.e., feature 1 (pattern (100)) mainly models the seven personality disorders (PDs), feature 2 (pattern (010)) models alcohol and drug abuse disorders and the antisocial PD, while feature 3 (pattern (001)) models anxiety and mood disorders. Additionally, Figure 3b shows the probability of being male and female for the latent feature vectors \mathbf{z}_n depicted in the legend and the empirical probability of being male and female in the database (baseline).

In Figure 3b, we observe that having no active features (pattern (000)), which captures people that do not suffer from any disorder, increases the probability of being male with respect to the baseline probability, and therefore, it indicates that females tend to suffer in a higher extent from psychiatric disorders. Additionally, we observe that feature 1 (pattern (100)) increases the probability of being male, while feature 3 (pattern (001)) increases the probability of being female. Hence, from the analysis of Figure 3b, we can conclude that, while women suffer more frequently from mood and anxiety disorders than men, PDs are more common in men.

4. Conclusions

In this paper, we have introduced the first available general latent feature model and its code implementation, which will ease researchers from diverse fields to analyze a wide range of heterogeneous, incomplete and noisy datasets in

an automatic manner. We have showed the flexibility and applicability of the proposed GLFM by performing data exploratory analysis of diverse real-world datasets.

References

- Blanco, C., Krueger, R. F., Hasin, D. S., Liu, S. M., Wang, S., Kerridge, B. T., Saha, T., and Olfson, M. Mapping common psychiatric disorders: Structure and predictive validity in the National Epidemiologic Survey on Alcohol and Related Conditions. *Journal of the American Medical Association Psychiatry*, 70(2):199–208, 2013.
- Byar, D. P. and Green, S. B. The choice of treatment for cancer patients based on covariate information: application to prostate cancer. *Bulletin du Cancer*, 67:477–490, 1980.
- Doshi-Velez, F. and Ghahramani, Z. Accelerated sampling for the indian buffet process. In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09*, pp. 273–280, New York, NY, USA, 2009. ACM.
- Gopalan, P., Ruiz, F. J. R., Ranganath, R., and Blei, D. M. Bayesian Nonparametric Poisson Factorization for Recommendation Systems. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2014.
- Griffiths, T. L. and Ghahramani, Z. The Indian buffet process: an introduction and review. *Journal of Machine Learning Research*, 12:1185–1224, 2011.
- Li, X.-B. A Bayesian approach for estimating and replacing missing categorical data. *J. Data and Information Quality*, 1(1): 3:1–3:11, June 2009. ISSN 1936-1955.
- Lunn, M. and McNeil, D. Applying cox regression to competing risks. *Biometrics*, 51(2):524–532, 1995.
- Ruiz, F. J. R., Valera, I., Blanco, C., and Perez-Cruz, F. Bayesian nonparametric modeling of suicide attempts. *Advances in Neural Information Processing Systems*, 25:1862–1870, 2012.
- Ruiz, F. J. R., Valera, I., Blanco, C., and Perez-Cruz, F. Bayesian nonparametric comorbidity analysis of psychiatric disorders. *Journal of Machine Learning Research*, 2013.
- Todeschini, A., Caron, F., and Chavent, M. Probabilistic low-rank matrix completion with adaptive spectral regularization algorithms. In Burges, C.J.C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K.Q. (eds.), *Advances in Neural Information Processing Systems 26*, pp. 845–853. Curran Associates, Inc., Dec. 2013.
- Valera, I., Fernandez-Pradier, M., and Ghahramani, Z. General Latent Feature Models for Heterogeneous Datasets. *ArXiv: https://arxiv.org/abs/1706.03779*, 2017.
- Weich, S. and Lewis, G. Poverty, unemployment, and common mental disorders: population based cohort study. *BMJ*, 317 (7151):115–119, 1998.
- Weissman, M. M., R., Bland, Joyce, P. R., Newman, S., Wells, J.E., and Wittchen, H.-U. Sex differences in rates of depression: cross-national perspectives. *Journal of Affective Disorders*, 29:77 – 84, 1993. Special Issue Toward a New Psychobiology of Depression in Women.






Figure 2. **Data exploration of a prostate cancer clinical trial.** We depict the effect of each latent feature on each attribute. Panels (a)-(d) shows different indicators of the prostate cancer, as well as the dose level of DES. Panel (d) corresponds to Prognosis Status, which indicates whether the patient either is alive or dies from one of the following three causes: vascular disease, prostatic cancer, or other reason. The baseline refers to the empirical distribution of each attribute in the whole dataset.

General Latent Feature Model



(a) Probability of suffering from each disorder



(b) Gender distribution

Figure 3. Feature effects including gender in the analysis. (a) Probabilities of suffering from the 20 considered disorders and (b) probability of being male and female for the latent feature vectors z_n shown in the legend and for the baseline.