SPARSE THREE-PARAMETER RESTRICTED INDIAN BUFFET PROCESS FOR UNDERSTANDING INTERNATIONAL TRADE

Melanie F. Pradier¹, Viktor Stojkoski², Zoran Utkovski², Lujupco Kocarev², and Fernando Perez-Cruz¹

¹ University Carlos III in Madrid; ² Macedonian Academy of Sciences and Arts

ABSTRACT

This paper presents a Bayesian nonparametric latent feature model specially suitable for exploratory analysis of high-dimensional count data. We perform a non-negative doubly sparse matrix factorization that has two main advantages: not only we are able to better approximate the row input distributions, but the inferred topics are also easier to interpret. By combining the three-parameter and restricted Indian buffet processes into a single prior, we increase the model flexibility, allowing for a full spectrum of sparse solutions in the latent space. We demonstrate the usefulness of our approach in the analysis of countries' economic structure. Compared to other approaches, empirical results show our model's ability to give easy-tointerpret information and better capture the underlying sparsity structure of data.

Index Terms— Bayesian nonparametrics, count data, infinite matrix factorization

1. INTRODUCTION

Exploration in high-dimensional data needs to balance predictive accuracy with interpretability [1]. When collaborating with experts in other fields, the primary goal is often not only reducing some error measure, but rather understanding the structure of data [2, 3, 4]. The data exploration phase can then be turned into a data exploitation phase via policy recommendations, medical protocols or as a further improved discriminative model [5].

Data exploration comes in different forms. PCA and factor analysis are linear methods that provide non-sparse solutions with strong Gaussianity assumptions. Local linear embedding [6], isomap [7] and Gaussian process latent variable models [8] learn non-linear manifolds in high dimensional spaces with non-sparse features; non-negative matrix factorization [9] provides a low dimensional sparse representation of the data. Also, Bayesian nonparametric (BNP) models can be used for clustering [10] and sparse feature analysis [11], in which the underlying latent dimension is unknown. Many applications have benefited from



Fig. 1: Data considered in this paper. A non-zero entry reflects a relative advantage of a country at exporting a given product. Note the triangular structure.

sparse data exploration models, like computer vision [12], genomics [13, 14], psychiatry [16], or sports [17].

This paper presents a non-negative matrix factorization model specially suitable for high-dimensional count data. Our model provides easy-to-interpret features and captures sparsity structure in data. This is illustrated in the context of international trade and economic growth [18, 19, 20]. Data considered in this paper is approximately triangular after reordering of rows and columns, as shown in Fig. 1: here, countries have different diversity degrees in their export portfolios, and thus different trade strategies and skills. Our objective is then to capture such triangular structure, e.g., discover the underlying capabilities of countries, and their relationships.

We rely on two different extensions of the Indian buffet process (IBP) to learn a potentially infinite number of latent features. The three-parameter IBP allows for different degrees of sharing between features [21], whereas the restricted IBP allows for a general distribution over the number of active features per row [22]. We combine both elements into a sparse Poisson factorization scheme. In the trade context, our model is able to represent different kinds of realities, from a world in which countries with few skills focus on different types of products, to a world in which poor countries have a strong overlap in export skills. We also allow for data points to exhibit different number of active features, e.g., we expect poor countries to have very little features active, while developed countries might have almost all features active. Our contributions include a novel BNP model for count data, a corresponding inference algorithm, and extensive empirical validation in the context of trade data.

Thanks to the European Union 7th Framework Programme through the Marie Curie ITN "Machine Learning for Personalized Medicine" MLPM2012 for funding.

2. THEORETICAL BACKGROUND

2.1. Indian buffet process (IBP)

The IBP is a stochastic process defining a probability distribution over equivalence classes of sparse binary matrices with a finite number of rows and unbounded number of columns [11]. Although the number of columns is potentially infinite, only a finite number of those will contain non-zero entries due to the finite nature of the observed data. The IBP can be derived taking the limit as $K \to \infty$ of a finite binary matrix $\mathbf{Z} \in \{0, 1\}^{N \times K}$, where N is the number of observations, and K is the number of latent features. Each element z_{nk} is generated as:

$$\pi_k \sim \text{Beta}(\alpha/K, 1),$$

 $z_{nk} \sim \text{Bernoulli}(\pi_k)$ (1)

where π_k is the probability of observing a non-zero value in column k, $\mathbf{Z}_{n\bullet}$ is the *n*-th row for sample *n* and $\mathbf{Z}_{\bullet k}$ is the *k*-th column for feature *k*. We say that a feature *k* is active for sample *n* if $z_{nk} = 1$. When $K \to \infty$, the above finite model tends to the IBP, denoted by: $\mathbf{Z} \sim \text{IBP}(\alpha)$, where α is the mass parameter controlling the a priori activation probability of new features.

In the IBP, the expected number of active features per row is distributed according to $\text{Poisson}(\alpha)$ and the total number of active features K^+ , i.e., number of columns with non-zero entries, is distributed as $\text{Poisson}(\alpha \sum_{i=1}^{N} \left(\frac{1}{i}\right))$. The single scalar parameter α has thus an effect on both the density (total number of ones) and sparsity structure (position of the non-zero values within **Z**). Such assumption might be too restrictive in general data exploration tasks, failing to capture situations such as a high number of latent features with low activation levels, or varying degrees of per-row sparsity in the latent matrix.

2.2. Extensions of the IBP

Three-Parameter IBP (3P-IBP) As its name indicates, the 3P-IBP can be fully specified by three parameters [21]: α is the same mass parameter from the IBP, $\sigma \in [0, 1)$ controls the power-law behavior of the model (weight decay), and $c > -\sigma$ is the concentration parameter that affects the *a priori* number of ones per column (sharing degree across features). When c = 1 and $\sigma = 0$, we recover the standard IBP model.

By introducing parameters c and σ , the latent matrix **Z** (defined in 2.1) has a more flexible sparsity structure, regardless of the sparsity density which is controlled by α . The 3P-IBP gives more flexibility on the feature weights, but has the disadvantage that the number of ones per-row is still Poisson distributed a priori for all data points, which might not be desirable in all scenarios. This problem can be directly addressed by the R-IBP.

Restricted IBP (R-IBP) The recently developed R-IBP allows for an arbitrary prior distribution f over the number of active features per row [22]. The R-IBP has two degrees of freedom α and f to respectively control for sparsity degree and sparsity structure of **Z**. The intuition behind the R-IBP is easy to explain in the commonly used culinary metaphor for IBPs, where rows design customers, and columns refer to dishes in an Indian buffet [11]. Customers in the R-IBP have varying degrees of hunger: some of them sample from many dishes in the buffet (the non-zero values), while others only taste a reduced set of dishes. This is generally convenient to model structured data, e.g., international trade, where developed countries are known to have more assets, and thus are expected to exhibit a higher number of latent features (capabilities) compared to poor countries.

3. OUR APPROACH

Modeling Let $\mathbf{X} \in \mathbb{N}^{N \times D}$ be our input matrix of N data points and D dimensions. We build an infinite latent feature model for count data with Poisson likelihood and Gamma-distributed factors:

$$x_{nd} \sim \operatorname{Poisson}(\mathbf{Z}_{n \bullet} \mathbf{B}_{\bullet d}),$$
 (2)

$$B_{kd} \sim \text{Gamma}\left(\alpha_B, \frac{\mu_B}{\alpha_B}\right),$$
 (3)

where \mathbf{Z} is a binary matrix, and α_B and μ_B are the shape and mean parameters of the prior Gamma distribution for each element of matrix \mathbf{B} . Sparsity in matrix \mathbf{B} can be induced simply by choosing $\alpha_B \ll 1$. Both \mathbf{Z} and \mathbf{B} are then non-negative and sparse, which makes the inferred latent variables particularly interpretable. To decouple sparsity density and sparsity structure in \mathbf{Z} , we combine the advantages of both the R-IBP and 3P-IBP into a single prior,

$$\mathbf{Z} \sim 3 \text{R-IBP}(\alpha, c, \sigma, f),$$
 (4)

where α, c, σ , and f refer to the parameters defined in Sec. 2.2. We rely on a negative binomial distribution for f, which is best understood as an overdispersed Poisson. Hence it will naturally allow for countries to exhibit a much variable range of active features. We refer to the model as *Sparse Three-parameter Restricted Indian buffet process* (S3R-IBP). It can be seen as a probabilistic extension of non-negative matrix factorization where the number of latent features is not fixed a priori, both matrices are sparse, and soft-constraints on the latent sparsity structure are imposed through a more flexible prior.

In the trade context, we have N countries, D products and K^+ non-empty latent features to be inferred. A given row $\mathbf{Z}_{n\bullet}$ captures which latent features (skills) are active for country n. Matrix **B** represents the effect of each latent feature on every product. For instance, if a latent **Algorithm 1** A single iteration of the MCMC inference procedure for the S3R-IBP model.

- 1: Sample each element of **Z** using inclusion probabilities [25, 22].
- 2: Sample latent measure π using MH steps [22].
- Sample each element of B and X from their conditional distributions (conjugate priors).
- 4: Sample hyperparameter α according to [10].

feature k is active for a certain country, all products having high values in vector $\mathbf{B}_{k\bullet}$ will be more likely to be exported by that country.

Inference Since exact computation of the posterior distribution for the latent variables is intractable, we resort to a Markov Chain Monte Carlo (MCMC) approach. Our algorithm uses Gibbs sampling together with Metropolis-Hasting (MH). Following [23], we introduce the auxiliary variables $x'_{nd,1}, \ldots, x'_{nd,K}$ for each observation x_{nd} such that $x_{nd} = \sum_{k=1}^{K} x'_{nd,k}$, and $x'_{nd,k} \sim \text{Poisson}(Z_{nk}B_{kd})$ for $k = 1, \ldots, K$. Given such auxiliary variables, the model is conditionally conjugate, and a Gibbs sampler can be derived straightforwardly. The complete sampling algorithm is described in Alg. 1.

4. EXPERIMENTS

Two publicly available trade datasets, the SITC and HS, are considered for the year 2010. The data represents the Revealed Comparative Advantage (RCA) of countries, a normalized common measure in economics [24] already illustrated in Fig. 1. Simulations are run for 10 different train-test splits with a proportion of 90-10% entries. The MCMC burn-in period is 30,000 iterations, and results are averaged using the last 1,000 posterior samples.

Model hyperparameters We choose f = Negative-Binomial(r, p), with r = [1, 2], and p = [0.1, 0.3, 0.5]. Results were equivalent using any of these priors. Here, we report results for r = 1 and p = 0.1. We also ran experiments for each combination of c = [1, 10, 20, 50] and $\sigma = [0, 0.25, 0.5, 0.75, 1]$. Parameter c was found to be more influential than σ . We here report the best setting c = 50, and $\sigma = 1$. Hyperparameters for the Gamma prior over α are shape and scale equal to one. Finally, α_B is set to 0.01 to induce sparsity, and $\mu_B = 1$.

4.1. Quantitative evaluation

Table 1 compares our model against probabilistic matrix factorization (MF) [26], non-negative MF (NMF) [27], IBP [11], and sparse IBP (S-IBP) which uses $\alpha_B < 1$, in terms of predictive accuracy and interpretability strength.

Accuracy All models present similar perplexity (the lower, the better), except S-IBP, in which the sparseness restriction degrades its performance significantly. S3R-IBP has the same sparsity constraint, but its more flexible prior compensates the penalty in perplexity, leading to a performance close to the non-sparse models, i.e. MF and IBP. The S3R-IBP match the perplexity performance of non sparse methods, but keeping the results interpretable.

Interpretability To assess semantic quality, we rely on coherence [28], which is an often-used metric in topic modeling literature. The closer coherence is to zero, the better. S3R-IBP outperforms IBP and S-IBP by far, making it specially suitable for data exploration in highdimensional count scenarios. The non-sparse methods present a very low coherence, as expected.

Sparsity structure Figure 2 evaluates the ability of S3R-IBP to fit the input distribution of the number of nonzero values per-row in **X**, versus IBP, S-IBP, and a simple binomial model from the economic literature [29]. We measure the "proximity" of the empirical and predicted distribution via qq-plots. S-IBP underfits the distribution for higher values, e.g., it predicts a lower number of countries with high number of exports, in contrast to the S3R-IBP model.

4.2. Qualitative evaluation

Interpretability Table 2 lists all the averaged latent features learned by S3R-IBP. Table 3 shows feature example F1 learned by IBP and S-IBP. Features are matched across the 10-folds using the Jaccard index similarity. S3R-IBP is able to give much shorter and concise descriptions, as weights decrease at a faster pace and reach higher values at the top. Products in the IBP list are



Fig. 2: Capturing sparsity structure. S3R-IBP gives the best fit for the distribution of number of non-zero values per row in X.

Metric	Metric MF		IBP	S-IBP	S3R-IBP			
Log Perplexity Coherence	$\begin{array}{c} 1.68 \pm 0.01 \\ -264.60 \pm 4.74 \end{array}$	$\begin{array}{c} 1.61 \pm 0.01 \\ -263.27 \pm 7.45 \end{array}$	$\begin{array}{c} {\bf 1.59 \pm 0.04} \\ -149.36 \pm 7.56 \end{array}$	3.26 ± 0.17 -178.44 ± 4.50	$\begin{array}{c} 1.62 \pm 0.01 \\ - 140.51 \pm 2.73 \end{array}$			
(a) 2010 SITC database ($N = 126$, $D = 744$, 16k non-zero values, 17% sparsity)								
Metric	MF	NMF	IBP	S-IBP	S3R-IBP			
Log Perplexity Coherence	$\begin{array}{c} 1.48 \pm 0.01 \\ -264.73 \pm 3.11 \end{array}$	$\begin{array}{c} {\bf 1.47 \pm 0.01} \\ {-264.67 \pm 6.22} \end{array}$	$\begin{array}{c} 1.58 \pm 0.01 \\ -148.91 \pm 10.57 \end{array}$	2.56 ± 0.12 -168.39 \pm 13.16	$\begin{array}{c} 1.57 \pm 0.02 \\ - 134.51 \pm 4.43 \end{array}$			

(b) 2010 HS database (N = 123, D = 4890, 77k non-zero values, 13% sparsity)

Table 1: **Quantitative evaluation of accuracy and interpretability.** S3R-IBP beats MF, NMF, IBP, and S-IBP in terms of topic coherence while retaining similar predictive accuracy (in terms of test log-perplexities).

Id	Products with highest weights	IBP
$\mathbf{F1}$	misc. animal oils (0.78) , bovine entails (0.72) , bovine meat (0.68) , milk (0.63) , equine (0.62) , butter (0.58)	confectionary sug
F2	synthetic woven, synth. yarn, woven $< 85\%$ synth.	baked goods
F3	parts metalworking, tool parts, polishing stones	tissue paper (
F4	Aldehyde–Ketone, glycosides–vaccines, medicaments	metal container
F5	synthetic rubber, acrylic polymers, silicones	soaps (0.3
F6	measuring instruments, math inst., electrical inst.	1
F7	vehicles parts, cars, iron wire	C IDD
F8	improved wood, mineral wool, heating equipment	S-IBP
F9	elect. machinery, vehicles stereos, data processing eq.	
F10	baked goods, metal containers, misc. edibles	bovine (0.
F11	misc. articles of iron, carpentry wood, wood articles	improved wood
F12	vegetables, fruit-vegetable juices, misc. fruit	misc. vegetable c
F13	misc. pumps, ash-residues, chemical wood pulp	butter (0.3
F14	synth. undergarments, feminine outerwear, men's shirts	rape seeds (
F15	mise rotating electric plant parts control inst of gas	misc. wheat

Table 2: Features learned by S3R-IBP.Products with higher weights are reported.

heterogeneous. The S-IBP list includes items from a mixture of farming and technological elements, whereas the S3R-IBP list is more homogeneous.

Features correlation To analyze the existing correlation between latent features, we apply our S3R-IBP on Z as input data. Such a deep structure, i.e., using a two-layer IBP, has already been explored in [30]. S3R-IBP infers two meta-features, M-F0 and M-F1, which assign different weights to each latent feature from the first layer. Countries with an active M-F1 are those that have more active features in the first layer and larger GDP. M-F1 can be interpreted as the meta-feature that distinguishes between developed countries and developing ones, resulting in a sharp division of the world in terms of capabilities.

M-F0 and M-F1 divide the original features into three disjoint sets. The first set contains features whose weight is either zero or insignificant in M-F1 (highlighted in red). These define countries with least capabilities, dealing with less complex products like farming or textile (see Table 2). The second set is composed by F10 (in green), which has a high value in both meta-features. This feature is present in both developing and developed countries, although developed countries do trade them more efficiently than developing ones (higher weights). The last set includes

IBP	Id	Weight	Id	Weight	
confectionary sugar (0.45) plastic containers (0.43) baked goods (0.41) tissue paper (0.40) metal containers (0.39) soaps (0.39)	F14 F12 F10 F2 F1 F9	$\begin{array}{c} 0.37 \\ 0.32 \\ 0.17 \\ 0.16 \\ 0.14 \\ 0.13 \end{array}$	F8 F11 F15 F10 F7 F6	$0.69 \\ 0.68 \\ 0.60 \\ 0.59 \\ 0.52 \\ 0.34$	
	F13	0.05	F13	0.32	
S-IBP	F6 F5	$0.04 \\ 0.04$	F4 F3	0.31	
bovine (0.53) improved wood (0.51) misc. vegetable oils (0.50)	F4 F15 F7	0.04 0.04 0.03	F5 F1 F9	0.14 0.05 0.02	
butter (0.50) rape seeds (0.47) misc. wheat (0.45)	F8 F11 F3	0.03 0.02 0.02	F2 F14 F12	$0.01 \\ 0.00 \\ 0.00$	
	(a) M-F0		(b	(b) M-F1	

Table 3: Competitors.Table 4: Meta-features.AExample matched to F1. sharp division of the world arises.

features whose weights in M-F0 are negligible compared to their weights in M-F1 (in black). These features contain products like chemicals and complex machinery, which are mostly traded by developed countries.

Such sharp division among features suggests the existence of a "poverty" or "quiescence trap" in the spirit of [29], a trap of development stasis in which some countries get stuck due to the inability to "acquire" capabilities associated with the production of more complex products.

5. CONCLUSION

This paper presents the S3R-IBP model, a non-negative and sparse infinite matrix factorization for data exploration of high-dimensional count data. The model is able to capture complex sparsity structure, and delivers compact, easy to interpret features. We illustrate the usefulness of this model to explain trade data, by interpreting latent features as country capabilities which are required for producing products. The presented approach is general for any other count-data scenario where broad, flexible assumptions are needed. As future work, we plan to introduce a Markovian dependency in the latent space, allowing for dynamic feature activation over time.

6. REFERENCES

- M. I. Jordan, "Message from the president: The era of big data," *ISBA Bulletin*, vol. 18, no. 2, pp. 1–3, 2011.
- [2] F. Doshi-Velez and B. Kim, "Towards A Rigorous Science of Interpretable Machine Learning," Feb. 2017, arXiv: 1702.08608.
- [3] B. Kim, J. A. Shah, and F. Doshi-Velez, "Mind the Gap: A Generative Approach to Interpretable Feature Selection and Extraction," in *Advances in Neural Information Processing Systems 28*, pp. 2260–2268. Curran Associates, Inc., 2015.
- [4] A. Vellido, J. D. Martín-Guerrero, and Paulo J.G. Lisboa, "Making machine learning models interpretable.," in *ESANN*. 2012, vol. 12, pp. 163–172, Citeseer.
- [5] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin, ""Why Should I Trust You?": Explaining the Predictions of Any Classifier," in *Proceedings of the 22Nd ACM SIGKDD*, New York, NY, USA, 2016, KDD '16, pp. 1135–1144, ACM.
- [6] S. T. Roweis and L. K. Saul, "Nonlinear Dimensionality Reduction by Locally Linear Embedding," *Science*, vol. 290, no. 5500, pp. 2323–2326, Dec. 2000.
- [7] J. B. Tenenbaum, V. Silva, and J. C. Langford, "A Global Geometric Framework for Nonlinear Dimensionality Reduction," *Science*, vol. 290, no. 5500, pp. 2319–2323, Dec. 2000.
- [8] N. Lawrence, "Probabilistic Non-linear Principal Component Analysis with Gaussian Process Latent Variable Models," *Journal of Machine Learning Research*, vol. 6, no. Nov, pp. 1783–1816, 2005.
- [9] P. O. Hoyer, "Non-negative Matrix Factorization with Sparseness Constraints," J. Mach. Learn. Res., vol. 5, pp. 1457–1469, Dec. 2004.
- [10] C. Archambeau, B. Lakshminarayanan, and G. Bouchard, "Latent IBP compound dirichlet allocation," *IEEE transactions on Pattern Analysis and Machine Intelligence*, vol. 37, no. 2, pp. 321–333, 2015.
- [11] T. L. Griffiths and Z. Ghahramani, "The Indian Buffet Process: An Introduction and Review.," *Journal of Machine Learning Research*, vol. 12, pp. 1185–1224, 2011.
- [12] J. Wright, Y. Ma, J. Mairal, G. Sapiro, T. S. Huang, and S. Yan, "Sparse Representation for Computer Vision and Pattern Recognition," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 1031–1044, June 2010.
- [13] D. Knowles and Z. Ghahramani, "Nonparametric Bayesian sparse factor models with application to gene expression modeling," *The Annals of Applied Statistics*, vol. 5, no. 2B, pp. 1534–1552, June 2011.
- [14] M. Chen, C. Gao, and H. Zhao, "Phylogenetic Indian Buffet Process: Theory and Applications in Integrative Analysis of Cancer Genomics," arXiv:1307.8229 [math, q-bio, stat], July 2013, arXiv: 1307.8229.

- [15] K. Miller, M. I. Jordan, and T. L. Griffiths, "Nonparametric Latent Feature Models for Link Prediction," in Advances in Neural Information Processing Systems 22, pp. 1276–1284. Curran Associates, Inc., 2009.
- [16] F. J. R. Ruiz, I. Valera, C. Blanco, and F. Perez-Cruz, "Bayesian Nonparametric Comorbidity Analysis of Psychiatric Disorders," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1215–1247, Jan. 2014.
- [17] M. F. Pradier, F. J. R. Ruiz, and F. Perez-Cruz, "Prior Design for Dependent Dirichlet Processes: An Application to Marathon Modeling," *PLOS ONE*, vol. 11, no. 1, pp. e0147402, Jan. 2016.
- [18] C. A. Hidalgo, B. Klinger, A.-L. Barabási, and R. Hausmann, "The Product Space Conditions the Development of Nations," *Science*, vol. 317, no. 5837, pp. 482–487, July 2007.
- [19] R. Hausmann, C. A. Hidalgo, S. Bustos, M. Coscia, S. Chung, J. Jimenez, A. Simoes, and M. A. Yildirim, "The Atlas of Economic Complexity," 2014.
- [20] V. Stojkoski, Z. Utkovski, and L. Kocarev, "The Impact of Services on Economic Complexity: Service Sophistication as Route for Economic Growth," 2016, arXiv: 1604.06284.
- [21] Y. W. Teh and D. Gorur, "Indian Buffet Processes with Power-law Behavior," in Advances in Neural Information Processing Systems 22, pp. 1838–1846. Curran Associates, Inc., 2009.
- [22] F. Doshi-Velez and S. A. Williamson, "Restricted Indian Buffet Processes," Aug. 2015, arXiv: 1508.06303.
- [23] P. Gopalan, J. M. Hofman, and D. M. Blei, "Scalable Recommendation with Poisson Factorization," Nov. 2013, arXiv: 1311.1704.
- [24] B. Balassa, "The Purchasing-Power Parity Doctrine: A Reappraisal," *Journal of Political Economy*, vol. 72, no. 6, pp. 584–596, 1964.
- [25] N. Aires, "Algorithms to find exact inclusion probabilities for conditional Poisson sampling and Pareto sampling designs," *Methodology and Computing in Applied Probability*, vol. 1, no. 4, pp. 457–469, 1999.
- [26] R. Salakhutdinov and A. Mnih, "Probabilistic matrix factorization," in Advances in Neural Information Processing Systems. 2008.
- [27] M. N. Schmidt, O. Winther, and L. K. Hansen, "Bayesian non-negative matrix factorization," in *Independent Com*ponent Analysis and Signal Separation, p. 540?547. 2009.
- [28] A. Fan, F. Doshi-Velez, and L. Miratrix, "Promoting Domain-Specific Terms in Topic Models with Informative Priors," arXiv:1701.03227 [cs], Jan. 2017.
- [29] R. Hausmann and C. A. Hidalgo, "The network structure of economic output," *Journal of Economic Growth*, vol. 16, no. 4, pp. 309–342, Oct. 2011.
- [30] F. Doshi-Velez and Z. Ghahramani, "Correlated nonparametric latent feature models," in *Proceedings of* the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence. 2009, pp. 143–150, AUAI Press.