
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Leppäaho, Eemeli; Ammad-ud-din, Muhammad; Kaski, Samuel

GFA

Published in:
Journal of Machine Learning Research

Published: 01/01/2017

Document Version
Publisher's PDF, also known as Version of record

Published under the following license:
CC BY

Please cite the original version:
Leppäaho, E., Ammad-ud-din, M., & Kaski, S. (2017). GFA: Exploratory Analysis of Multiple Data Sources with Group Factor Analysis. *Journal of Machine Learning Research*, 18, 1-5. [39]. <http://jmlr.org/papers/v18/16-509.html>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

GFA: Exploratory Analysis of Multiple Data Sources with Group Factor Analysis

Eemeli Leppäaho

EEMELI.LEPPAAHO@AALTO.FI

Muhammad Ammad-ud-din

MUHAMMAD.AMMAD-UD-DIN@AALTO.FI

Samuel Kaski

SAMUEL.KASKI@AALTO.FI

Helsinki Institute for Information Technology HIIT

Department of Computer Science

Aalto University

P.O.Box 15400, FI-00076 Aalto, FINLAND

Editor: Mohammad Emtiyaz Khan

Abstract

The R package GFA provides a full pipeline for factor analysis of multiple data sources that are represented as matrices with co-occurring samples. It allows learning dependencies between subsets of the data sources, decomposed into latent factors. The package also implements sparse priors for the factorization, providing interpretable biclusters of the multi-source data.

Keywords: Bayesian latent variable modelling, biclustering, data integration, factor analysis, multi-view learning

1. Introduction

The need for intelligent data integration and analysis methods is getting increasingly common. We often encounter connected measurements from multiple data sources, possessing complex shared variation and structured noise. Modelling the relationships within this kind of a data collection, and using them for prediction, is a challenging problem. Two existing R packages, CCAGFA (Klami et al., 2015) and CMF (Klami et al., 2014), provide good tools for this¹, but remain limited in their use for two reasons: the packages do not provide a complete data analysis pipeline, and they assume the structure present within data sources to be dense, that is, present in all the samples and the features. This assumption is often unrealistic, and we provide package GFA to allow sparse factorization of multi-source data, along with tools from data preprocessing to model interpretation.

Extending from the classical Bayesian matrix factorization work, *group factor analysis* (GFA) (Virtanen et al., 2012; Klami et al., 2015) assumes that data matrix $\mathbf{Y}^{(m)} \in \mathbb{R}^{N \times D_m}$ can be approximated with $\mathbf{X}\mathbf{W}^{(m)\top}$, where the latent variable matrix $\mathbf{X} \in \mathbb{R}^{N \times K}$ describes the observed data in a lower dimensional space, and the projection matrix $\mathbf{W}^{(m)} \in \mathbb{R}^{D_m \times K}$ gives a mapping between the latent space and the data space of the m th source. A key aspect of the model definition is a group sparse prior for the \mathbf{W} matrices: the association

1. Available at <https://cran.r-project.org/package=CCAGFA> and <https://cran.r-project.org/package=CMF>.

of a data matrix to component k is based on evidence in the data. Allowing only a subset of the components to explain a data matrix results in the posterior identifying relationships between any subset of the data sources. The spike-and-slab sparsity priors used to achieve this are described by Bunte et al. (2016). Alternative priors aiming for similar goals have been presented by Zhao et al. (2016)², though not containing any tools for preprocessing or model interpretation.

2. The GFA Package

GFA implements a Gibbs sampling scheme for inferring the posterior of group factor analysis, and it is available for R at <https://cran.r-project.org/package=GFA>. It has been designed to contain the key elements needed for the data analysis task of factorization of multiple data sources.

2.1 Model Options and Initialization

The package includes a variety of model options for the user to choose from, depending on the application requirement. The default priors assume latent factors that are dense along the sample space, and group sparsity along the feature space (Klami et al., 2015). Using them, the posterior given data Y can be inferred by:

```
opts <- getDefaultOpts()
res <- gfa(Y,opts)
```

Another main use of the package is biclustering of multiple data sources, which is done by using element-wise sparse priors for the matrices \mathbf{X} and \mathbf{W} . This approach has been presented in Bunte et al. (2016), and is usable via `getDefaultOpts(bicluster=TRUE)`.

The input data Y given to function `gfa()` is a list of matrices (that is, data sources) that each have an equal number of rows. Alternatively, if there are several data sources with the same columns, it is possible to input a list with two elements: the sources with paired rows, and the sources with paired columns. Other important options include which parts of the data collection share the scale of each component or the strength of the residual noise (by default each data source for both), and whether the convergence of the sampling chain should be estimated (based on the Geweke diagnostic of its reconstruction). A full description of the possible options is given in the documentation of the function `getDefaultOpts()`.

2.2 Missing Value Prediction

GFA can be used for prediction by simply including the missing matrix elements as `NA`s. The model parameters are sampled based on the observed data only, and the missing values are predicted with the function `reconstruction()`. If the data samples are acquired in a sequential manner, it is alternatively possible to i) infer the models parameters on the first batch, and ii) using fixed projections, quickly predict missing data sources of the second batch based on the observed data sources. Both the approaches use learned relationships between the data sources (shared components) to predict from one to other, and provide predictive distributions along with point estimates.

2. Available at <https://github.com/judyboon/BASS>.

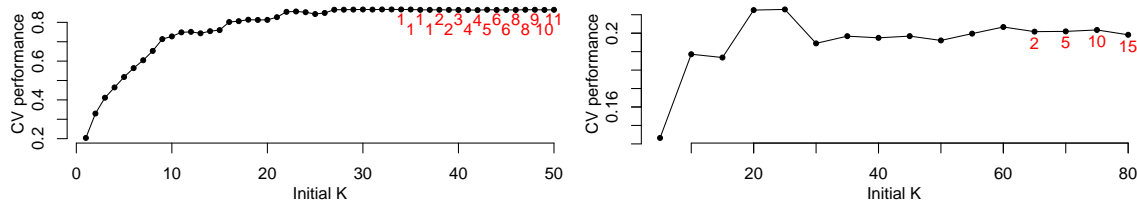


Figure 1: Illustrating model complexity selection with GFA on synthetic data (left) and GDSC data set (right), showing the cross-validated predictive performance (Spearman correlation), and the number of empty components (red numbers) as a function of initial model complexity K . High initial K allows the components to produce a flexible noise model, as opposed to low initial K .

3. Usage

In addition to inferring model parameters and predicting missing values, the data analysis pipeline typically needs additional steps presented here.

3.1 Preprocessing

The input data can be normalized in different ways with `normalizeData(Y)`. As GFA aims to explain maximal variance with a compact set of parameters, and has no separate assumptions for the bias (mean) of features, centering the data is advised. Furthermore, the higher the variance of features (or data sources), the larger their effect to the model parameters. If the user considers this to be desired behavior, no scaling should be done. On the other hand, if the features (or data sources) are deemed approximately equally important, they should be scaled accordingly.

3.2 Model Complexity Selection

The priors of GFA are able to infer the model complexity. This is achieved by initializing the model with a high number of components K (default $\min(N, \sum_m D_m)/2$) and allowing the sparsity prior to prune out excessive components (explaining only little variance). As complexity selection is a difficult task in general, we recommend using any prior information of the data, if available, when selecting the initial K . Here it should be noted that the components can model (structured) noise in addition to the signal of interest, inducing a flexible noise model. If this is not desired, `informativeNoisePrior()` can be used to specify how large proportion of variance is expected to be explained with the components; in effect causing weaker (structured noise) components to be explained by the simple residual noise. The package issues a warning if the automatic model complexity selection does not seem to work as intended, for example if no components are pruned.

We demonstrate model complexity selection with two examples in Figure 1: a synthetic data set, and a cancer cell line data set by Genomics of Drug Sensitivity in Cancer project (GDSC) (Yang et al., 2013). The complexity of the synthetic data set is overestimated slightly, resulting in accurate predictions as long as initial K is above the true $K = 32$. With the GDSC data set, where the true model complexity is unknown, initial K at 20 or 25 gives the best predictions, but the full model is not far behind.

3.3 Interpreting the Model

The inferred factorization explains dependencies in the data collection through K components. The factorization can be visualized using function `visualizeComponents()`, which shows the component activities in the data sources. Shared components imply predictability from a data source to another. The function can also show the full parameter matrices \mathbf{X} and \mathbf{W} , as well as their effect, and the effect of individual components, in the data space.

3.4 Robustness Analysis

As in factor analysis in general, the factorization GFA aims to infer is unidentifiable, resulting in an extremely multi-modal posterior.³ Thus consideration is needed if a user wishes to analyze multiple factorizations (in this case sampling chains) jointly, as they are not likely to be identical. To the best of our knowledge, there are no sampling methods that would in practice converge to the true multi-modal posterior with high-dimensional data. To this end, we provide function `robustComponents()` which uses a correlation-based procedure to analyze which components occur frequently in different GFA sampling chains. This approach can be used to obtain robust factorizations for noisy data collections, where the components explain only small parts of the variation, resulting in uncertain estimates.

3.5 Examples

The package manual contains simplified examples demonstrating function usage. Besides this, we provide the following three demonstrations:

- `demo(GFAPipeline)`: Simple illustration of the GFA pipeline.
- `demo(GFAexample)`: Elaborate illustration of the GFA model on simulated data.
- `demo(GFAdream)`: Replication of Bunte et al. (2016) results (requiring data download, which is instructed).

4. Discussion

GFA allows inferring the relationships between multiple co-occurring data matrices by reconstructing them with group sparse priors for the projection matrices. It finds a decomposition with components that are specific to the data sources, shared between them all, or shared between any subset of them. This kind of sparsity is beneficial for exploratory data analysis and integration. Our GFA package covers essential tools ranging from preprocessing to modelling assumptions and from robustness analysis to interpreting the model.

Acknowledgments

This work was financially supported by the Academy of Finland (Finnish Center of Excellence in Computational Inference Research COIN; grants 295503 and 292337 to MA and SK). We acknowledge the computational resources provided by Aalto Science-IT project.

3. Unimodal posterior could be achieved in special cases, given limiting identifiability constraints.

References

- Kerstin Bunte, Eemeli Leppäaho, Inka Saarinen, and Samuel Kaski. Sparse group factor analysis for biclustering of multiple data sources. *Bioinformatics*, 32(16):2457–2463, 2016. Preprint: <http://arxiv.org/abs/1512.08808>.
- Arto Klami, Guillaume Bouchard, and Abhishek Tripathi. Group-sparse embeddings in collective matrix factorization. In *International Conference on Learning Representations*, 2014.
- Arto Klami, Seppo Virtanen, Eemeli Leppäaho, and Samuel Kaski. Group factor analysis. *IEEE Transactions on Neural Networks and Learning Systems*, 26(9):2136–2147, 2015.
- Seppo Virtanen, Arto Klami, Suleiman A. Khan, and Samuel Kaski. Bayesian group factor analysis. In Neil Lawrence and Mark Girolami, editors, *Fifteenth International Conference on Artificial Intelligence and Statistics*, volume 22 of *JMLR W&CP*, pages 1269–1277. JMLR, 2012.
- Wanjuan Yang, Jorge Soares, Patricia Greninger, Elena J Edelman, Howard Lightfoot, Simon Forbes, Nidhi Bindal, Dave Beare, James A Smith, I Richard Thompson, et al. Genomics of drug sensitivity in cancer (GDSC): a resource for therapeutic biomarker discovery in cancer cells. *Nucleic Acids Research*, 41(D1):D955–D961, 2013.
- Shiwen Zhao, Chuan Gao, Sayan Mukherjee, and Barbara E Engelhardt. Bayesian group factor analysis with structured sparsity. *Journal of Machine Learning Research*, 17(196): 1–47, 2016.