



# StepMix: A Python Package for Pseudo-Likelihood Estimation of Generalized Mixture Models with External Variables

Sacha Morin 

Université de Montréal  
Mila - Quebec AI Institute

Robin Legault 

Massachusetts Institute of Technology  
Operations Research Center

Félix Laliberté 

Université de Montréal

Zsuzsa Bakk 

Leiden University

Charles-Édouard Giguère 

Institut universitaire en santé  
mentale de Montréal

Roxane de la Sablonnière 

Université de Montréal

Éric Lacourse 

Université de Montréal

---

## Abstract

**StepMix** is an open-source Python package for the pseudo-likelihood estimation (one-, two- and three-step approaches) of generalized finite mixture models (latent profile and latent class analysis) with external variables (covariates and distal outcomes). In many applications in social sciences, the main objective is not only to cluster individuals into latent classes, but also to use these classes to develop more complex statistical models. These models generally divide into a measurement model that relates the latent classes to observed indicators, and a structural model that relates covariates and outcome variables to the latent classes. The measurement and structural models can be estimated jointly using the so-called one-step approach or sequentially using stepwise methods, which present significant advantages for practitioners regarding the interpretability of the estimated latent classes. In addition to the one-step approach, **StepMix** implements the most important stepwise estimation methods from the literature, including the bias-adjusted three-step methods with Bolk-Croon-Hagenaars and maximum likelihood corrections and the more recent two-step approach. These pseudo-likelihood estimators are presented in this paper under a unified framework as specific expectation-maximization subroutines. To facilitate and promote their adoption among the data science community, **StepMix** follows the object-oriented design of the **scikit-learn** library and provides an additional R wrapper.

*Keywords:* mixture models, expectation-maximization, Python, R.

---

## 1. Introduction

Mixture models are a family of probabilistic models that can be estimated from observed data to discover hidden or latent subgroups within a population. They are often used to analyze multivariate continuous and categorical data following some explicit assumptions about their conditional distribution given a categorical latent variable. Mixture modeling belongs to the field of model-based cluster analysis (McLachlan, Lee, and Rathnayake 2019) and takes different names across academic disciplines (Sterba 2013). In health and social sciences, it is often referred to as latent class analysis when the observed variables are categorical and latent profile analysis when they are continuous (Oberski 2016). Mixture models can also be presented under the more general framework of probabilistic graphical models as directed acyclic graphs with latent variables (Koller and Friedman 2009).

In social sciences, mixture models are used in cross-sectional and longitudinal studies that require to account for discrete population heterogeneity. For example, they are employed in psychology, psychiatry and epidemiology to establish probabilistic diagnoses when a gold standard is unavailable. Using national survey data, Lacourse, Baillargeon, Dupéré, Vitaro, Romano, and Tremblay (2010) identified subgroups of adolescents that are more likely to present specific diagnoses of conduct disorder (non-aggressive vs. aggressive). Other applications based on cross-sectional data include identifying patterns of mobile internet usage for traveling (Okazaki, Campo, Andreu, and Romero 2015) and modeling recidivism rates among latent classes of juvenile offenders (Mulder, Vermunt, Brand, Bullens, and Van Merle 2012).

Historically, two main approaches have been used for estimating the parameters of mixture models. While it can be seen as a single likelihood maximization problem (Dayton and Macready 1988), the estimation process can also be divided into distinct steps. The first stepwise estimators of regression models became popular during the 1960s (Goldberger 1961). In order to relate latent classes to observed indicators and distal outcomes, these methods predict the class membership of the units based on the observed indicator variables before estimating the class-conditional distribution of the outcomes using the predicted latent classes. Their main advantage over classical maximum likelihood estimators (MLE) lies in the fact that they avoid the distal outcomes to contribute to the definition of the latent variables. This property is essential in applications where the interpretability of the latent classes is an important consideration. Among other advantages, stepwise methods are useful in the presence of missing data on distal outcomes, as the unobserved outcome variables only affect the last step of the estimation procedure. Unfortunately, several studies highlighted that naive stepwise methods produce highly biased estimates of the class-conditional distribution of the distal outcomes (Croon 2002; Bolck, Croon, and Hagenaars 2004; Vermunt 2010; Devlieger, Mayer, and Rosseel 2016). To alleviate this issue, stepwise estimators that mitigate or completely eliminate the bias introduced in the classification process have been proposed in the literature (Bolck *et al.* 2004; Vermunt 2010; Bakk and Kuha 2018).

Although commercial packages implement modern bias-adjusted stepwise estimators, these methods still have very limited availability in open-source software. Furthermore, they are completely unavailable in Python (Van Rossum and Drake 2009), which considerably reduces their pool of potential users. As the popularity of mixture modeling increases among applied researchers, so does the importance of making the state-of-the-art estimation methods publicly available in a well-documented software package. The goal of **StepMix** is to respond to this need by providing the scientific community with a convenient open-source

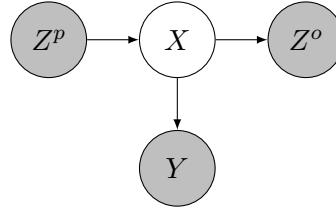


Figure 1: Family of mixture models that can be estimated by **StepMix**. Conditioning on the latent class  $X$  blocks the paths between the observed indicators, illustrating graphically the conditional independence assumptions and how the joint likelihood in (1) factorizes.

package implementing bias-adjusted stepwise estimators. Our Python package is distributed via the Python Package Index (PyPI) at <https://pypi.org/project/stepmix/>. An R wrapper is also available on the Comprehensive R Archive Network (CRAN) at <https://CRAN.R-project.org/package=stepmixr>.

This paper aims to introduce the **StepMix** package and present a self-contained reference on pseudo-likelihood estimation of mixture models with external variables. The family of mixture models studied in this work is formally defined in Section 2. Section 3 presents an overview of the estimation methods developed in the literature for these models and of the software packages in which they are implemented. Section 4 provides a detailed exposition of the estimation methods implemented in **StepMix** under a unified framework. The basic usage of our package and some of its most important features are discussed in Section 5. Computational examples based on simulated and real-life datasets are presented in Section 6. Section 7 concludes the paper.

## 2. Models

This section introduces the structure of the mixture models **StepMix** can estimate and the main notation used in the paper. In its most general formulation, the model we consider is composed of a latent categorical variable  $X$ , a set of observed indicators  $Y$ , a set of observed covariates (or predictors)  $Z^p$ , and a set of observed distal outcomes  $Z^o$ . The support of each of these sets of variables is respectively denoted by  $\mathcal{X}$ ,  $\mathcal{Y}$ ,  $\mathcal{Z}^p$ , and  $\mathcal{Z}^o$ . As illustrated in Figure 1, the observed variables  $Z^p$ ,  $Y$ , and  $Z^o$  are assumed to be conditionally independent given the latent class  $X$ .

Throughout the paper, for the sake of notational simplicity, the name of the variables are omitted in the formulas. For example,  $p(x, y)$  denotes the joint probability  $p(X=x, Y=y)$ . Furthermore, the notation  $p(\cdot)$  indistinctly refers to the probability mass function (PMF) in the case of discrete variables and to the probability density function (PDF) when handling continuous variables.

In accordance with the literature, we will distinguish the measurement model (MM) from the structural model (SM). They are respectively composed of the subsets of variables  $(X, Y)$  and  $(Z^p, X, Z^o)$ . The MM is specified by a set of parameters  $\theta_M \in \Theta_M$  that can be partitioned as  $\theta_M = (\theta_X, \theta_Y) \in (\Theta_X \times \Theta_Y)$ , where  $\theta_X$  and  $\theta_Y$  respectively specify the marginal distribution of the latent classes and the class-conditional distribution of the indicator variables. The MM

defines, for any realization  $(x, y) \in \mathcal{X} \times \mathcal{Y}$ , the joint probability:

$$p(x, y; \theta_M) = p(x; \theta_X)p(y | x; \theta_Y)$$

Analogously, we denote the parameters of the SM by  $\theta_S = (\theta_{Z^p}, \theta_{Z^o})$ , where the sets of parameters  $\theta_{Z^p}$  and  $\theta_{Z^o}$  respectively specify the distribution of the latent classes conditionally on the observed covariates and the class-conditional distribution of the distal outcomes. In the SM, the joint probability of a realization  $(x, z^o) \in \mathcal{X} \times \mathcal{Z}^o$  given the covariates  $z^p \in \mathcal{Z}^p$  is:

$$p(x, z^o | z^p; \theta_S) = p(x | z^p; \theta_{Z^p})p(z^o | x; \theta_{Z^o})$$

In the absence of covariates, the latent class is treated as an exogenous variable in the SM. In this case, the SM's parameters reduce to  $\theta_S = \theta_{Z^o}$  and only specify the class-conditional distribution of the distal outcomes.

Together, the MM and the SM form the complete model (CM), which specifies the probabilities:

$$p(x, y, z^o | z^p; \theta_M, \theta_S) = p(x | z^p; \theta_X, \theta_{Z^p})p(y | x; \theta_Y)p(z^o | x; \theta_{Z^o}) \quad (1)$$

We adopt a generative perspective on the CM when it does not contain covariates. This means that the marginal distribution of all the variables, including the latent classes, is explicitly specified by the model. In particular, given fixed parameters  $(\theta_X, \theta_Y, \theta_{Z^o})$ , this allows to directly sample observations  $(x, y, z^o)$  from the CM. On the other hand, a conditional perspective is adopted in the presence of covariates. In other words, the covariates  $Z^p$  are regarded as exogenous variables, and their marginal distribution is not estimated. Observations  $(x, y, z^o)$  can then be sampled from the CM given fixed covariates  $z^p$  and parameters  $(\theta_Y, \theta_{Z^p}, \theta_{Z^o})$ . In this case, the parameters  $\theta_X$  specifying the marginal distribution of the latent classes are ignored in the CM. The conditional distribution of the latent variable in (1) thus reduces to:

$$p(x | z^p; \theta_X, \theta_{Z^p}) = \begin{cases} p(x; \theta_X) & \text{if } Z^p = \emptyset, \\ p(x | z^p; \theta_{Z^p}) & \text{otherwise.} \end{cases} \quad (2)$$

### 3. Related packages

Two classical approaches have been and are still widely applied in the literature for estimating the mixture models illustrated in Figure 1. The first is the one-step approach, which directly estimates the parameters of the CM through likelihood maximization. This estimation is typically performed using the expectation-maximization (EM) algorithm (Dempster, Laird, and Rubin 1977). The second is the so-called naive three-step approach. Its first step consists of computing the MLE of the MM's parameters without considering the SM. In the second step, the latent class of each unit is predicted based on the estimated MM. In step three, after replacing the latent classes with their predicted values from step two, the parameters of the SM are estimated by likelihood maximization. The one-step and naive three-step approaches are based on standard EM, classification, and maximum likelihood estimation procedures. They can thus be implemented relatively easily based on any software for latent class modeling. Nevertheless, the commercial software packages **Latent GOLD** (Vermunt and Magidson 2013) and **Mplus** (Muthén and Muthén 2017) are the only ones that currently support three-step estimators natively.

The first bias-adjusted stepwise estimator that was introduced in the literature is the Bolk-Croon-Hagenaars three-step method (BCH), named after its authors (Bolck *et al.* 2004). The original BCH approach was improved by Vermunt (2010), who also proposed the so-called ML three-step method. The main idea of bias-adjusted methods is to estimate the probability of misclassifying units in the second step in order to perform a correction of the imputed class weights used in step three. These methods are implemented in **Latent GOLD** and **Mplus**, but are unavailable in open-source software.

An alternative stepwise estimator is the two-step approach (Bakk and Kuha 2018). This method bypasses the classification step of the three-step approach and thus eliminates the bias resulting from misclassification errors. It follows the logic of stepwise regression modeling (Goldberger 1961) by first estimating the parameters of the MM before performing a pseudo-maximum likelihood estimation of the CM in which the estimated MM's parameters are held fixed. The two-step approach is available in **Latent GOLD** and **Mplus**, and can easily be implemented in any software package allowing for conditioning on a set of fixed parameters. Unfortunately, most model-based clustering open-source packages do not allow using fixed parameters in this manner. The only exception is the recent R (R Core Team 2025) package **multilevLCA** (Lyrvall, Mari, Bakk, Oser, and Kuha 2023).

These stepwise estimation methods all fall within the frequentist paradigm of pseudo-maximum likelihood estimation. Although this is outside of the scope of this paper, it is worth mentioning that a Bayesian perspective on mixture modeling can also be adopted. In particular, maximum a posteriori (MAP) estimation makes it possible to include prior information in the analysis and is exploited as a regularization technique in model-based clustering (Fraley and Raftery 2007). In addition to MAP estimation based on the EM algorithm, **BayesLCA** (White and Murphy 2014) supports Gibbs sampling (Geman and Geman 1984) and a variational Bayes approximation method (Ormerod and Wand 2010). The packages **depmixS4** (Visser and Speekenbrink 2010), **randomLCA** (Beath 2017), **mclust** (Scrucca, Fop, Murphy, and Raftery 2016) and **scikit-learn** (Pedregosa *et al.* 2011) also implement Bayesian methods in the context of one-step estimation, primarily for regularization purposes.

A list of popular R and Python packages for the estimation of mixture models follows.

- **mclust** (Scrucca *et al.* 2016) is an R package implementing Gaussian mixture models with different covariance structures. The package also features functions for model-based hierarchical clustering and model selection. **AutoGMM** (Athey, Liu, Pedigo, and Vogelstein 2019) is a Python package with similar features.
- **MoEClust** (Murphy and Murphy 2020) is an R package for parsimonious finite multivariate Gaussian mixtures of experts models. The package supports **mclust**'s different covariance structures and also gating/expert network covariates.
- **scikit-learn** (Pedregosa *et al.* 2011) offers a number of Gaussian mixture models in Python and is a **StepMix** dependency.
- **multilevLCA** (Lyrvall *et al.* 2023) is an R package for multilevel latent class models. It allows for the inclusion of covariates using the one-step and two-step estimators but does not support distal outcomes. It is currently the only other open-source package that implements a stepwise approach.

Package	Version	R	Python	scikit-learn API	Two-step estimation	Bias-adjusted three-step estimation	Gaussian and non-Gaussian components	Covariates
<b>StepMix</b>	2.1.3	✓	✓	✓	✓	✓	✓	✓
<b>scikit-learn</b>	1.2		✓	✓				
<b>multilevLCA</b>	1.1	✓			✓		✓	✓
<b>mclust</b>	6.0	✓						
<b>MoEClust</b>	1.5.2	✓						✓
<b>AutoGMM</b>	2.0.1		✓	✓				
<b>MixMod</b>	0.2.0		✓				✓	
<b>Rmixmod</b>	2.1.8	✓					✓	
<b>poLCA</b>	1.6.0.1	✓						✓
<b>depmixS4</b>	1.5	✓					✓	✓
<b>randomLCA</b>	1.1-2	✓						
<b>BayesLCA</b>	1.9	✓						
<b>e1071::lca</b>	1.7-13	✓						
<b>glca</b>	1.3.3	✓						✓
<b>VarSelLCM</b>	2.1.3.1	✓				✓		
<b>FlexMix</b>	2.3-18	✓				✓	✓	

Table 1: Features of open-source software packages for mixture models estimation.

- **FlexMix** (Grün and Leisch 2008) is an R package that provides a flexible framework for finite mixture models of generalized linear models. It supports models with covariates as well as several families of probability distributions.
- **PoLCA** (Linzer and Lewis 2011) is an R package for one-step estimation of latent class models with categorical observed variables and covariates. It does not support Gaussian components.
- **MixMod** and **Rmixmod** (Lebret, Iovleff, Langrognet, Biernacki, Celeux, and Govaert 2015) are the Python and R interfaces to the C++ (Stroustrup 2013) **mixmodLib** library. These packages are based on a generalized linear model perspective and implement model-based clustering methods for quantitative and qualitative data. They provide different variants of the EM algorithm, including the stochastic EM and the classification EM (see McLachlan and Krishnan 2007 for a presentation of these methods). However, these packages do not support models with covariates.

In Table 3, the main features of **StepMix** are compared with those of the previously mentioned software as well as the other open-source R and Python packages, LCA from **e1071** (Meyer *et al.* 2024), **glca** (Kim, Jeon, Chang, and Chung 2022), and **VarSelLCM** (Marbac and Sedki 2019). Table 3 indicates the programming languages of each package and, for Python packages, whether they follow a **scikit-learn** API. It also reports which packages natively support two-step and three-step estimation methods in addition to the standard one-step approach. The last two columns indicate whether each package supports both Gaussian and non-Gaussian indicators and outcome variables and whether the SM can contain covariates.

**StepMix** is the first open-source software that natively implements bias-adjusted three-step methods. In addition to enabling the users of Python to use mixture models with covariates and both Gaussian and non-Gaussian components, **StepMix** is also the first package to support

pseudo-likelihood estimation of mixture models in this language. One of our package's primary advantages over open-source and commercial existing software is that it strictly follows the **scikit-learn** interface. All the tools implemented in **scikit-learn**, including model selection and hyperparameter tuning methods, can thus be directly applied to Python's **StepMix** estimator. In social sciences, **Latent GOLD** and **Mplus** are still the main tools used to estimate latent class models, mostly because of the lack of comprehensive packages in Python or R for mixture modeling. These software packages implement one-step, two-step, and bias-adjusted three-step estimators, with some limitations regarding the available distributions for non-categorical outcome variables. However, their code implementation of estimation methods is not open-source and thus cannot be directly inspected by researchers. Furthermore, these packages are expensive, making them inaccessible to a large part of the scientific community, hence the still frequent use of the one-step and naive three-step approaches in the literature. By bridging the gap between the features of commercial and open-source software **StepMix**, aims to foster the adoption of the state-of-the-art estimation methods of mixture models among applied researchers.

## 4. Pseudo-likelihood estimation methods

The main objective of the estimation methods implemented in **StepMix** is to approximate the parameters of the generalized mixture model illustrated in Figure 1 using the maximum likelihood principle. In the estimation process, the MM and the SM can be considered either jointly or separately. This section presents the stepwise estimation methods from the literature under a unified framework in which each approach is based on the maximization of specific log-likelihood and pseudo-log-likelihood functions using the EM algorithm. The log-likelihood functions of the MM, SM, and CM are first presented in Section 4.1. Then, the EM algorithm is derived for the CM in Section 4.2. In Sections 4.3 to 4.5, the one-step, two-step, and three-step methods implemented in **StepMix** are described as specific EM subroutines that are applied on the MM, SM, and CM. This presentation highlights the fact that the stepwise estimation methods from the literature all share a very similar structure and primarily differ by the parameters that are considered fixed in the estimation process and, for three-step methods, by the definition of imputed class weights.

### 4.1. Likelihood functions

Suppose we have a sample of units  $i \in N$ , each with observed covariates, indicators, and distal outcomes  $(Z_i^p, Y_i, Z_i^o) = (z_i^p, y_i, z_i^o)$ . Also, let the support of the latent variables  $X_i$  be given by a set of  $K$  classes  $\mathcal{X} = \{1, \dots, K\}$ . Finally, suppose that each unit has been given a sample weight  $\omega_i \in \mathbb{R}$ . The log-likelihood function of the MM, SM, and CM are then respectively denoted by  $\ell_N^M$ ,  $\ell_N^S$  and  $\ell_N^C$ , and are given by:

$$\begin{aligned}\ell_N^M(\theta_M \mid y) &= \sum_{i \in N} \omega_i \log p(y_i; \theta_M) \\ &= \sum_{i \in N} \omega_i \log \left( \sum_{k \in \mathcal{X}} p(X_i=k, y_i; \theta_M) \right) \\ &= \sum_{i \in N} \omega_i \log \left( \sum_{k \in \mathcal{X}} p(X_i=k; \theta_X) p(y_i \mid X_i=k; \theta_Y) \right)\end{aligned}$$

$$\begin{aligned}
\ell_N^S(\theta_S \mid z^p, z^o) &= \sum_{i \in N} \omega_i \log p(z_i^o \mid z_i^p; \theta_S) \\
&= \sum_{i \in N} \omega_i \log \left( \sum_{k \in \mathcal{X}} p(X_i=k, z_i^o \mid z_i^p; \theta_S) \right) \\
&= \sum_{i \in N} \omega_i \log \left( \sum_{k \in \mathcal{X}} p(X_i=k \mid z_i^p; \theta_{Z^p}) p(z_i^o \mid X_i=k; \theta_{Z^o}) \right) \\
\ell_N^C(\theta_M, \theta_S \mid y, z^o \mid z^p) &= \sum_{i \in N} \omega_i \log p(y_i, z_i^o \mid z_i^p; \theta_M, \theta_S) \\
&= \sum_{i \in N} \omega_i \log \left( \sum_{k \in \mathcal{X}} p(X_i=k, y_i, z_i^o \mid z_i^p; \theta_M, \theta_S) \right) \\
&= \sum_{i \in N} \omega_i \log \left( \sum_{k \in \mathcal{X}} p(X_i=k \mid z_i^p; \theta_X, \theta_{Z^p}) \right. \\
&\quad \left. p(y_i \mid X_i=k; \theta_Y) p(z_i^o \mid X_i=k; \theta_{Z^o}) \right)
\end{aligned}$$

**StepMix** supports missing data in the indicators and distal outcomes through full information maximum likelihood (FIML). This method was first outlined by [Hartley and Hocking \(1971\)](#), and is a direct generalization of maximum likelihood estimation to datasets containing data points with missing observed variables. The FIML approach makes it possible to use all available information despite missing data by defining the probabilities in the likelihood function only with respect to the observed variables for each unit. Formally, let us consider a MM with  $D$  independent indicators. Suppose that the indicators  $y_{id}$  of unit  $i \in N$  are missing for each dimension  $d \in \bar{\mathcal{D}}_i \subseteq \mathcal{D} = \{1, \dots, D\}$ . In the FIML framework, the probability  $p(y_i \mid X_i=k; \theta_Y) = \prod_{d \in \mathcal{D}} p(y_{id} \mid X_i=k; \theta_Y)$  is replaced by  $p(y_i \mid X_i=k; \theta_Y) = \prod_{d \in \mathcal{D} \setminus \bar{\mathcal{D}}_i} p(y_{id} \mid X_i=k; \theta_Y)$ . If no indicator is observed, i.e.,  $\bar{\mathcal{D}}_i = \mathcal{D}$ , then we fix  $p(y_i \mid X_i=k; \theta_Y) = 1$ . In the absence of missing data, FIML reduces to classical likelihood maximization. Everywhere in the paper, including in the definition of the above log-likelihood functions, the probabilities  $p(y_i \mid X_i=k; \theta_Y)$  and  $p(z_i^o \mid X_i=k; \theta_{Z^o})$  are therefore assumed to respect the definition required for the use of FIML.

## 4.2. EM algorithm

The log-likelihood functions of Section 4.1 are generally multi-modal and non-convex due to the presence of a sum within the logarithm. The computation of their MLE is thus not straightforward and is tackled via the EM algorithm in **StepMix**.

The EM algorithm, introduced by [Dempster et al. \(1977\)](#), is the standard approach in the literature for maximum likelihood estimation in statistical models with latent categorical variables. The main idea of this method is to replace the log probability of each observation by a lower bounding auxiliary function to obtain a pseudo-likelihood function that is easier to maximize than the original log-likelihood function.

In the following, we derive the EM algorithm for the CM. For each unit  $i \in N$ , the log probability of the observations  $(y_i, z_i^o)$ , given the covariates  $z_i^p$  and the CM's parameters  $(\theta_M, \theta_S)$ , is underestimated by an auxiliary function  $\mathcal{L}_i(q_i, \theta_M, \theta_S)$ , which is parameterized by a probability distribution  $q_i(\cdot)$ .

$$\begin{aligned} \log p(y_i, z_i^o | z_i^p; \theta_M, \theta_S) &= \log \left( \sum_{x \in \mathcal{X}} p(x, y_i, z_i^o | z_i^p; \theta_M, \theta_S) \right) \\ &= \log \left( \sum_{x \in \mathcal{X}} \frac{q_i(x)p(x, y_i, z_i^o | z_i^p; \theta_M, \theta_S)}{q_i(x)} \right) \end{aligned} \quad (3)$$

$$\begin{aligned} &= \log \left( \mathbb{E}_{q_i} \left[ \frac{p(x, y_i, z_i^o | z_i^p; \theta_M, \theta_S)}{q_i(x)} \right] \right) \\ &\geq \mathbb{E}_{q_i} \left[ \log \left( \frac{p(x, y_i, z_i^o | z_i^p; \theta_M, \theta_S)}{q_i(x)} \right) \right] \end{aligned} \quad (4)$$

$$=: \mathcal{L}_i(q_i, \theta_M, \theta_S)$$

For (3) to hold, the only condition on  $q(\cdot)$  is that  $q(x) > 0 \ \forall x \in \mathcal{X}$ . In equation 4 results from Jensen's inequality, using the fact that the logarithm function is concave.

From there, we want to select the probability distribution  $q_i(\cdot)$  that maximizes the lower bounding function  $\mathcal{L}_i(q_i, \theta_M, \theta_S)$  for it to approximate  $\log p(y_i, z_i^o | z_i^p; \theta_M, \theta_S)$  as closely as possible. To do so, we start by rewriting  $\mathcal{L}_i(q_i, \theta_M, \theta_S)$  in terms of the Kullback–Leibler (KL) divergence of distribution  $q(\cdot)$  from the conditional distribution of the latent variable  $X_i$  given the observations of unit  $i$  and the CM's parameters.

$$\begin{aligned} \mathcal{L}_i(q_i, \theta_M, \theta_S) &= \sum_{x \in \mathcal{X}} q_i(x) \log \left( \frac{p(x, y_i, z_i^o | z_i^p; \theta_M, \theta_S)}{q_i(x)} \right) \\ &= \sum_{x \in \mathcal{X}} q_i(x) \log \left( \frac{p(x | z_i^p, y_i, z_i^o; \theta_M, \theta_S) p(y_i, z_i^o | z_i^p; \theta_M, \theta_S)}{q_i(x)} \right) \\ &= \sum_{x \in \mathcal{X}} q_i(x) \left( \log \left( \frac{p(x | y_i, z_i^o | z_i^p; \theta_M, \theta_S)}{q_i(x)} \right) \right. \\ &\quad \left. + \log (p(y_i, z_i^o | z_i^p; \theta_M, \theta_S)) \right) \\ &= \sum_{x \in \mathcal{X}} q_i(x) \log \left( \frac{p(x | z_i^p, y_i, z_i^o; \theta_M, \theta_S)}{q_i(x)} \right) \\ &\quad + \sum_{x \in \mathcal{X}} q_i(x) \log (p(y_i, z_i^o | z_i^p; \theta_M, \theta_S)) \\ &= - \sum_{x \in \mathcal{X}} q_i(x) \log \left( \frac{q_i(x)}{p(x | z_i^p, y_i, z_i^o; \theta_M, \theta_S)} \right) \\ &\quad + \sum_{x \in \mathcal{X}} q_i(x) \log (p(y_i, z_i^o | z_i^p; \theta_M, \theta_S)) \\ &= -D_{\text{KL}}(q_i(\cdot) \| p(\cdot | z_i^p, y_i, z_i^o; \theta_M, \theta_S)) \\ &\quad + \sum_{x \in \mathcal{X}} q_i(x) \log (p(y_i, z_i^o | z_i^p; \theta_M, \theta_S)) \end{aligned}$$

$$\begin{aligned}
&= -D_{\text{KL}}(q_i(\cdot) \| p(\cdot | z_i^p, y_i, z_i^o; \theta_M, \theta_S)) \\
&\quad + \log(p(y_i, z_i^o | z_i^p; \theta_M, \theta_S)) \sum_{x \in \mathcal{X}} q_i(x) \\
&= -D_{\text{KL}}(q_i(\cdot) \| p(\cdot | z_i^p, y_i, z_i^o; \theta_M, \theta_S)) + \log(p(y_i, z_i^o | z_i^p; \theta_M, \theta_S)) \quad (5)
\end{aligned}$$

Maximizing (5) with respect to  $q(\cdot)$  over the set  $\mathcal{Q}$  of probability distributions that are strictly positive over  $\mathcal{X}$  is then relatively straightforward.

$$\begin{aligned}
\arg \max_{q_i \in \mathcal{Q}} \{\mathcal{L}_i(q_i, \theta_M, \theta_S)\} &= \arg \max_{q_i \in \mathcal{Q}} \{-D_{\text{KL}}(q_i(\cdot) \| p(\cdot | z_i^p, y_i, z_i^o; \theta_M, \theta_S)) \\
&\quad + \log(p(y_i, z_i^o | z_i^p; \theta_M, \theta_S))\} \\
&= \arg \max_{q_i \in \mathcal{Q}} \{-D_{\text{KL}}(q_i(\cdot) \| p(\cdot | z_i^p, y_i, z_i^o; \theta_M, \theta_S))\} \quad (6)
\end{aligned}$$

$$\begin{aligned}
&= \arg \min_{q_i \in \mathcal{Q}} \{D_{\text{KL}}(q_i(\cdot) \| p(\cdot | z_i^p, y_i, z_i^o; \theta_M, \theta_S))\} \\
&= p(\cdot | z_i^p, y_i, z_i^o; \theta_M, \theta_S) \quad (7)
\end{aligned}$$

(6) is obtained by removing the second term from the maximization problem. This can be done since this term does not depend on the distribution  $q_i(\cdot)$  with respect to which the expression must be maximized. From there, the problem reduces to minimizing a KL divergence. We use the fact that the distribution that minimizes the KL divergence from a distribution  $p(\cdot)$  is  $p(\cdot)$  itself to obtain the result in (7).

After fixing the probability distributions  $q_i(\cdot)$  for each unit  $i \in N$ , the log-likelihood function  $\ell_N^C(\theta_M, \theta_S | z^p, y, z^o)$  of the CM can be approximated by replacing the log probability of each unit  $i \in N$  by its auxiliary function  $\mathcal{L}_i(q_i, \theta_M, \theta_S)$ . The resulting pseudo-likelihood can then be maximized to approximate the MLEs of the original log-likelihood function.

$$\begin{aligned}
&\arg \max_{(\theta_M, \theta_S) \in (\Theta_M, \Theta_S)} \left\{ \sum_{i \in N} \omega_i \mathcal{L}_i(q_i, \theta_M, \theta_S) \right\} \\
&= \arg \max_{(\theta_M, \theta_S) \in (\Theta_M, \Theta_S)} \left\{ \sum_{i \in N} \omega_i \left( \sum_{x \in \mathcal{X}} q_i(x) \log(p(x, y_i, z_i^o | z_i^p; \theta_M, \theta_S)) \right. \right. \\
&\quad \left. \left. - \sum_{x \in \mathcal{X}} q_i(x) \log(q_i(x)) \right) \right\} \\
&= \arg \max_{(\theta_M, \theta_S) \in (\Theta_M, \Theta_S)} \left\{ \sum_{i \in N} \omega_i \left( \sum_{x \in \mathcal{X}} q_i(x) \log(p(x, y_i, z_i^o | z_i^p; \theta_M, \theta_S)) \right. \right. \\
&\quad \left. \left. - \sum_{i \in N} \omega_i \left( \sum_{x \in \mathcal{X}} q_i(x) \log(q_i(x)) \right) \right) \right\} \\
&= \arg \max_{(\theta_M, \theta_S) \in (\Theta_M, \Theta_S)} \left\{ \sum_{i \in N} \omega_i \left( \sum_{x \in \mathcal{X}} q_i(x) \log(p(x, y_i, z_i^o | z_i^p; \theta_M, \theta_S)) \right) \right\} \quad (8)
\end{aligned}$$

Analogously to (6), (8) is obtained by removing the terms that do not depend on the decision variables from the optimization problem. (8) corresponds to a likelihood maximization problem under complete information, as the latent variables  $X_i$  are now treated as observed variables. The parameters  $(\theta_M, \theta_S) \in (\Theta_M, \Theta_S)$  maximizing the objective function can thus be computed as in a standard MLE setting, without latent variables.

The EM algorithm starts from an initial set of estimated parameters  $(\hat{\theta}_M^{(0)}, \hat{\theta}_S^{(0)}) \in (\Theta_M \times \Theta_S)$ . At each iteration, the distributions  $q_i(x)$ ,  $i \in N$  are updated using the incumbent estimates of the CM parameters. More precisely, for each unit  $i \in N$  and each class  $k \in \mathcal{X}$ , we will denote the so-called class responsibilities by  $\tau_{i_k}^{(t)} := q_i^{(t)}(X=k) = p(X=k | z_i^p, y_i, z_i^o; \hat{\theta}_M^{(t)}, \hat{\theta}_S^{(t)})$ , which corresponds to the conditional distribution of  $X_i$  given the observations of unit  $i$  and the current estimated CM's parameters at iteration  $t \in \{0, 1, \dots\}$ . This step is called the expectation step (E Step) since it corresponds to maximizing the expectation in (4). The estimated CM's parameters are then updated by solving the complete information maximum likelihood problem given in (8) for the current imputed class weights  $q_i^{(t)}(X=k) = \tau_{i_k}^{(t)}$ . This latter step is referred to as the maximization step (M step).

This process is repeated iteratively until a user-defined convergence criterion is satisfied. As extensively discussed in [Wu \(1983\)](#), under mild conditions, the EM algorithm is guaranteed to converge to a local maximum of the original log-likelihood function. In practice, the EM algorithm is generally executed multiple times using different starting points to identify a global maximum with high probability. The convergence criteria implemented in **StepMix** are the relative and absolute gap between the average log-likelihood of subsequent estimated parameters.

### 4.3. One-step method

For a set of units  $i \in N$  with observed covariates, indicators, and distal outcomes  $(Z_i^p, Y_i, Z_i^o) = (y_i, z_i^p, z_i^o)$ , the objective of the one-step estimation method is to compute the MLE  $(\hat{\theta}_M^C, \hat{\theta}_S^C) = \arg \max_{(\theta_M, \theta_S) \in \Theta_M \times \Theta_S} \ell_N^C(\theta_M, \theta_S | z^p, y, z^o)$  of the CM's parameters. In **StepMix**, this is done by applying the EM algorithm on the CM directly, as presented in Algorithm 1. The notation  $(\hat{\theta}_M^C, \hat{\theta}_S^C)$  makes explicit the fact that the parameters  $\theta_M$  and  $\theta_S$  of the MM and SM are estimated by maximizing the log-likelihood function  $\ell_N^C(\cdot)$  of the CM in the one-step method. It intends to avoid notational ambiguity between this section and Sections 4.4 and 4.5. The lighter notation  $(\hat{\theta}_M, \hat{\theta}_S)$  is used in the pseudocode.

The one-step approach is the most natural way to estimate the CM's parameters based on a single dataset in which the observable variables of both the MM and the SM are available for all the units. In this setting, the one-step approach is generally the method that minimizes the bias and the variance of the CM's parameter estimates. This has been illustrated in several computational studies ([Vermunt 2010](#); [Asparouhov and Muthén 2014](#); [Bakk and Kuha 2018](#)), and is also visible in the results of Section 6. Nonetheless, as discussed in the same works, this method suffers from important practical and theoretical limitations. First, when using the one-step approach, model selection can be computationally cumbersome in the context of exploratory studies in which the number of potential indicator variables or covariates is large. Indeed, each time a variable is added to or removed from the model, both the SM and the MM must be re-estimated. This can cause computation time issues when using complex models and large datasets. The main theoretical drawback of this method, however, lies in the fact that it makes the estimators of the MM, and thus the interpretation of the latent classes, dependent on the SM. As a consequence, researchers cannot properly study the relation between a given latent variable and multiple research questions using the one-step approach. Finally, this method requires both the MM and the SM's parameters to be estimated using the same set of units. In practice, this can lead to methodological issues and limit the quality of the estimators of the MM's parameters. For example, a group of researchers may publish a MM model estimated on a very large sample or census data without being allowed to share

**Algorithm 1** One-step method

---

 Initialize  $(\hat{\theta}_M^{(0)}, \hat{\theta}_S^{(0)}) \in (\Theta_M \times \Theta_S)$ 


---

**STEP 1 : EM algorithm on the CM**Set  $t = 0$ **repeat**

$$\tau_{ik}^{(t)} = p(X_i=k|y_i, z_i^o; z_i^p, \hat{\theta}_M^{(t)}, \hat{\theta}_S^{(t)}), \forall i \in N, \forall k \in \mathcal{X} \quad \text{E Step} \quad \triangleright \text{Update responsibilities}$$

$$\hat{\theta}_M^{(t+1)} = \arg \max_{\theta_M \in \Theta_M} \sum_{i \in N} \omega_i \sum_{k \in \mathcal{X}} \tau_{ik}^{(t)} \log p(X_i=k, y_i; \theta_M) \quad \triangleright \text{Update } \hat{\theta}_M$$

$$\hat{\theta}_S^{(t+1)} = \arg \max_{\theta_S \in \Theta_S} \sum_{i \in N} \omega_i \sum_{k \in \mathcal{X}} \tau_{ik}^{(t)} \log p(X_i=k, z_i^o; z_i^p, \theta_S) \quad \triangleright \text{Update } \hat{\theta}_S$$

$$t = t + 1$$

**until** convergence

$$\text{Set } (\hat{\theta}_M, \hat{\theta}_S) = (\hat{\theta}_M^{(t)}, \hat{\theta}_S^{(t)}) \quad \triangleright \text{Estimated CM parameters}$$


---

$$\text{Return } (\hat{\theta}_M, \hat{\theta}_S) \quad \triangleright \text{Return estimated CM parameters}$$


---

the data, for example, for privacy reasons. In this case, a second group of researchers wishing to use the same MM would need to re-estimate it using their own dataset, possibly smaller, to estimate the parameters of their SM. The primary motivation of multi-step methods is thus to overcome these limitations by eliminating the dependency of the MM's parameters estimates on the SM.

#### 4.4. Two-step method

The two-step method requires a set of units  $i \in N_1$  with observed indicators  $Y_i = y_i$  and a set of units  $j \in N_2$  with observed covariates, indicators, and distal outcomes  $(Z_j^p, Y_j, Z_j^o) = (z_j^p, y_j, z_j^o)$ . There are no further constraints on samples  $N_1$  and  $N_2$ . In particular,  $N_1$  and  $N_2$  can be disjoint, but it is also possible that  $N_1 = N_2$  or  $N_2 \subset N_1$ .

This approach aims at computing the MLE  $\hat{\theta}_M^M = \arg \max_{\theta_M \in \Theta_M} \ell_{N_1}^M(\theta_M | y)$  of the MM's parameters on sample  $N_1$  and the maximum pseudo-likelihood estimator (MPLE)  $\hat{\theta}_S^{C|M} = \arg \max_{\theta_S \in \Theta_S} \ell_{N_2}^C(\theta_S | z^p, y, z^o; \hat{\theta}_M)$  of the CM's parameters on sample  $N_2$  conditionally on a fixed estimator  $\hat{\theta}_M$  of the MM model parameters.

In the first step, the EM algorithm is directly applied on the MM based on sample  $N_1$ . In the second step, the EM algorithm is applied on the CM for sample  $N_2$ , but the estimated MM's parameters  $\hat{\theta}_M$  obtained in the first step are held fixed instead of being re-estimated. This procedure is detailed in Algorithm 2.

The idea of the two-step approach was first suggested by [Bandeen-Roche, Miglioretti, Zeger, and Rathouz \(1997\)](#) as a method of evaluating the mutual sensitivity of the MM and SM's parameters in mixture models with discrete outcomes. It has then been applied by [Xue and Bandeen-Roche \(2002\)](#) in the context of SMs with covariates and by [Bartolucci, Montanari, and Pandolfi \(2014\)](#) in the case of latent Markov models for longitudinal data. Quite recently, [Bakk and Kuha \(2018\)](#) carried out the first comprehensive study of the two-step approach for generalized mixture models and provided empirical evidence for the superiority of the

**Algorithm 2** Two-step methodInitialize  $(\hat{\theta}_M^{(0)}, \hat{\theta}_S^{(0)}) \in (\Theta_M \times \Theta_S)$ **STEP 1 : EM algorithm on the MM**Set  $t = 0$ **repeat**

$$\tau_{ik}^{(t)} = p(X_i=k|y_i; \hat{\theta}_M^{(t)}), \forall i \in N_1, \forall k \in \mathcal{X} \quad \text{E Step} \quad \triangleright \text{Update responsibilities}$$

$$\hat{\theta}_M^{(t+1)} = \arg \max_{\theta_M \in \Theta_M} \sum_{i \in N_1} \omega_i \sum_{k \in \mathcal{X}} \tau_{ik}^{(t)} \log p(X_i=k, y_i; \theta_M) \quad \text{M Step} \quad \triangleright \text{Update } \hat{\theta}_M$$

$$t = t + 1$$

**until** convergence

$$\text{Set } \hat{\theta}_M = \hat{\theta}_M^{(t)} \quad \triangleright \text{Estimated MM parameters}$$

**STEP 2 : EM algorithm on the CM with  $\hat{\theta}_M$  fixed**Set  $t = 0$ **repeat**

$$\tau_{jk}^{(t)} = p(X_j=k|y_j, z_j^o, z_j^p; \hat{\theta}_M, \hat{\theta}_S^{(t)}), \forall j \in N_2, \forall k \in \mathcal{X} \quad \text{E Step} \quad \triangleright \text{Update responsibilities}$$

$$\hat{\theta}_S^{(t+1)} = \arg \max_{\theta_S \in \Theta_S} \sum_{j \in N_2} \omega_j \sum_{k \in \mathcal{X}} \tau_{jk}^{(t)} \log p(X_j=k, z_j^o, z_j^p; \theta_S) \quad \text{M Step} \quad \triangleright \text{Update } \hat{\theta}_S$$

$$t = t + 1$$

**until** convergence

$$\text{Set } \hat{\theta}_S = \hat{\theta}_S^{(t)} \quad \triangleright \text{Estimated SM parameters}$$

$$\text{Return } (\hat{\theta}_M, \hat{\theta}_S) \quad \triangleright \text{Return estimated CM parameters}$$

two-step approach over the widely used naive and bias-adjusted three-step methods in terms of bias and mean-squared error of the SM's parameters estimates. The two-step approach is an example of a two-stage pseudo-maximum likelihood estimation method (Besag 1975; Gong and Samaniego 1981), and thus inherits from the theoretical properties of this class of estimators. In particular, this implies that the two-step estimators of the SM's parameters are consistent. For most applications in which the joint estimation of the MM and the SM's parameters is undesirable, the two-step approach should be preferred to three-step methods, as it avoids introducing a classification error that would negatively affect the quality of the SM's parameters estimates computed in the last step.

#### 4.5. Three-step method

Like the two-step method, the three-step approach is based on a set of units  $i \in N_1$  with observed indicators  $Y_i = y_i$  and a set of units  $j \in N_2$  with observed covariates, indicators, and distal outcomes  $(Z_j^p, Y_j, Z_j^o) = (z_j^p, y_j, z_j^o)$ . Its first objective is also to compute the MLE  $\hat{\theta}_M^M = \arg \max_{\theta_M \in \Theta_M} \ell_{N_1}^M(\theta_M | y)$  of the MM's parameters based on sample  $N_1$ . From there, the goal is to compute a MPLE  $\hat{\theta}_S^S = \arg \max_{\theta_S \in \Theta_S} \tilde{\ell}_{N_2}^S(\theta_S | z; w)$  of the SM's parameters for sample  $N_2$ . The pseudo-log-likelihood function  $\tilde{\ell}_{N_2}^S(\theta_S | z; w)$  is defined in a different way for

each variant of the three-step method, but, in each case, only depends on the MM through a vector of imputed class weights  $w_j \in \mathbb{R}^K$  for each unit  $j \in N_2$ .

After applying the EM algorithm to the MM in the first step, imputed class weights are computed in the second step based on the estimated MM's parameters. In the naive version with soft assignment, the imputed class weights of each unit correspond to the class membership probabilities given by the MM. If modal assignment is used instead, each unit is completely assigned to its most likely latent class. The imputed class weights are then used in the third step to estimate the SM's parameters through a single iteration of the EM algorithm in which the responsibilities are fixed to the class weights obtained in the second step.

In the absence of uncertainty on the real latent classes, under usual regularity conditions on the distributions of the SM's variables, the naive three-step method would inherit the consistency and asymptotic normality properties of the MLE. However, the assignation phase of the three-step methods generally suffers from a non-negligible classification error rate which leads to biased estimates of the SM's parameters. In particular, [Bolck et al. \(2004\)](#) demonstrated that the uncorrected three-step approach underestimates the strength of the relationship between the distal outcomes variable and class membership. This can be intuitively understood by noticing that classification errors introduce units drawn from other classes in the subsample used to estimate the structural parameters of a given latent class. Thus, the higher the classification error rate, the more similar these subsamples become in terms of true class membership. As a consequence, the true conditional distribution of the outcomes given a class  $c \in \mathcal{X}$  impacts the estimated conditional distribution of the outcomes for other classes  $s \in \mathcal{X}$ . In practice, high misclassification rates thus lead to excessively similar estimated class-conditional distributions.

Two bias-adjusted three-step methods, namely the BCH ([Bolck et al. 2004](#); [Vermunt 2010](#)) and ML ([Vermunt 2010](#); [Bakk, Tekle, and Vermunt 2013](#)) methods, have been developed to alleviate the downward bias of the naive three-step approach. They both include in their second step the computation of a left stochastic matrix  $D \in \mathbb{R}^{K \times K}$  whose component  $(c, k)$  contains the probability  $p(W=k | X=c; \hat{\theta}_M)$  of assigning to class  $k$  a unit that actually belongs to class  $c$ , given the estimated MM's parameters  $\hat{\theta}_M$ . As explained by [Vermunt \(2010\)](#), if the support  $\mathcal{Y}$  of the indicator variables  $Y$  is finite, this probability can be computed exactly by summing over all the possible realizations of  $Y$ .

$$D_{ck} = p(W=k | X=c; \hat{\theta}_M) = \frac{\sum_{y \in \mathcal{Y}} p(X=c | Y=y; \hat{\theta}_M) \bar{w}_k^y}{p(X=c; \hat{\theta}_M)}, \quad \forall (c, k) \in \mathcal{X}^2$$

Note that the definition of matrix  $D$  depends on the type of assignment  $\mathcal{A} \in \{\text{soft, modal}\}$  used in the second step through the weight  $\bar{w}_k^y$  given to class  $k$  for a unit with observed indicators  $y$ . This value is defined as follows.

$$\bar{w}_k^y = \begin{cases} p(X=k | Y=y; \hat{\theta}_M) & , \text{ if } \mathcal{A} = \text{soft}, \\ \mathbf{1}[\arg \max_{c \in \{1, \dots, C\}} p(X | Y=y; \hat{\theta}_M) = k] & , \text{ if } \mathcal{A} = \text{modal}. \end{cases}$$

In the presence of continuous indicators, or if  $|\mathcal{Y}|$  is too large for the previous expression to be computed in a reasonable time, the probability  $p(W=k | X=c; \hat{\theta}_M)$  can be estimated using the empirical distribution given by the observed sample  $N_2$ . This latter definition is natively used in **StepMix** for computational efficiency concerns. Matrix  $D$  is thus computed as:

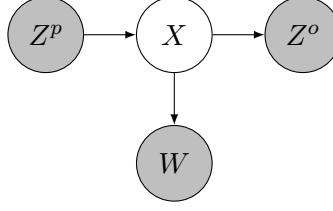


Figure 2: Model used in the third step of the three-step method with ML correction.

$$D_{ck} = \hat{p}(W=k \mid X=c; \hat{\theta}_M) = \frac{\sum_{j \in N_2} p(X=c \mid Y=y_j; \hat{\theta}_M) w_{jk}}{p(X=c; \hat{\theta}_M)}, \forall (c, k) \in \mathcal{X}^2$$

where  $w_{jk} = \bar{w}_k^{y_j}$ .

The third step of the BCH method is identical to that of the naive three-step method. It consists of estimating the SM's parameters through a single iteration of the EM algorithm using imputed class weights. The only difference is that, in BCH, these weights are computed based on the inverse of matrix  $D$  and, for each unit  $j \in N_2$ , are typically negative, except for the most probable class  $k$  based on the observed indicators  $y_j$  and the estimated MM's parameters  $\hat{\theta}_M$ . We refer the interested reader to [Vermunt \(2010\)](#); [Bakk \*et al.\* \(2013\)](#) for a detailed discussion on the BCH method.

In the case of the ML correction, the third step consists in splitting each unit  $j \in N_2$  into  $K$  copies  $k \in \mathcal{X}$  with weight  $w_{jk}$  and predicted class membership  $W_j=k$ , and applying the standard EM algorithm with multiple iterations on the model illustrated in Figure 2. This model is identical to the CM, except that the observed indicator variables are replaced by the predicted class membership, whose class-conditional distribution is given by  $p(W=k \mid X=c; \hat{\theta}_M) = D_{ck}$ .

At the  $t$ -th iteration of the EM algorithm, the responsibility of the  $k$ -th copy of unit  $j \in N_2$ , with predicted class membership  $W_j=k$ , is given by:

$$\begin{aligned}
\tau_{j_k c}^{(t)} &= p(X_j=c \mid z_j^p, W_j=k, z_j^o; \hat{\theta}_M, \hat{\theta}_S^{(t)}) \\
&= \frac{p(W_j=k, z_j^o \mid z_j^p, X_j=c; \hat{\theta}_M, \hat{\theta}_S^{(t)}) p(X_j=c \mid z_j^p; \hat{\theta}_M, \hat{\theta}_S^{(t)})}{p(W_j=k, z_j^o \mid z_j^p; \hat{\theta}_M, \hat{\theta}_S^{(t)})} \\
&= \frac{p(W_j=k \mid X_j=c; \hat{\theta}_M) p(z_j^o \mid z_j^p, X_j=c; \hat{\theta}_S^{(t)}) p(X_j=c \mid z_j^p; \hat{\theta}_M, \hat{\theta}_S^{(t)})}{p(W_j=k, z_j^o \mid z_j^p; \hat{\theta}_M, \hat{\theta}_S^{(t)})} \\
&= \frac{D_{ck} p(z_j^o \mid z_j^p, X_j=c; \hat{\theta}_S^{(t)}) p(X_j=c \mid z_j^p; \hat{\theta}_M, \hat{\theta}_S^{(t)})}{p(W_j=k, z_j^o \mid z_j^p; \hat{\theta}_M, \hat{\theta}_S^{(t)})} \\
&= D_{ck} \frac{p(z_j^o \mid z_j^p, X_j=c; \hat{\theta}_S^{(t)}) p(X_j=c \mid z_j^p; \hat{\theta}_M, \hat{\theta}_S^{(t)})}{p(z_j^o \mid z_j^p; \hat{\theta}_S^{(t)})} \frac{p(z_j^o \mid z_j^p; \hat{\theta}_S^{(t)})}{p(W_j=k, z_j^o \mid z_j^p; \hat{\theta}_M, \hat{\theta}_S^{(t)})} \\
&= D_{ck} p(X_j=c \mid z_j^p, z_j^o; \hat{\theta}_S^{(t)}) \frac{p(z_j^o \mid z_j^p; \hat{\theta}_S^{(t)})}{p(W_j=k, z_j^o \mid z_j^p; \hat{\theta}_M, \hat{\theta}_S^{(t)})}
\end{aligned}$$

$$\propto D_{ck} p(X_j=c \mid z_j^p, z_j^o; \hat{\theta}_S^{(t)}) p(z_j^o \mid z_j^p; \hat{\theta}_S^{(t)}) \quad (9)$$

In practice, to avoid having to handle a duplicated dataset in the third step of the three-step ML method, we express the responsibility  $\tau_{jc}^{(t)}$  of class  $c \in \mathcal{X}$  for each unit  $j \in N_2$  as the weighted sum of the responsibilities  $\tau_{jk}^{(t)}$ .

$$\begin{aligned} \tau_{jc}^{(t)} &= \sum_{k \in \mathcal{X}} w_{jk} \tau_{jk}^{(t)} \\ &= \sum_{k \in \mathcal{X}} w_{jk} D_{ck} p(X_j=c \mid z_j^p, z_j^o; \hat{\theta}_S^{(t)}) \frac{p(z_j^o \mid z_j^p; \hat{\theta}_S^{(t)})}{p(W_j=k, z_j^o \mid z_j^p; \hat{\theta}_M, \hat{\theta}_S^{(t)})} \end{aligned} \quad (10)$$

$$\lesssim \sum_{k \in \mathcal{X}} w_{jk} D_{ck} p(X_j=c \mid z_j^p, z_j^o; \hat{\theta}_S^{(t)}) p(z_j^o \mid z_j^p; \hat{\theta}_S^{(t)}) \quad (11)$$

$$\begin{aligned} &= p(z_j^o \mid z_j^p; \hat{\theta}_S^{(t)}) \sum_{k \in \mathcal{X}} w_{jk} D_{ck} p(X_j=c \mid z_j^p, z_j^o; \hat{\theta}_S^{(t)}) \\ &\propto \sum_{k \in \mathcal{X}} w_{jk} D_{ck} p(X_j=c \mid z_j^p, z_j^o; \hat{\theta}_S^{(t)}) \\ &= w_j D_c^T p(X_j=c \mid z_j^p, z_j^o; \hat{\theta}_S^{(t)}) \end{aligned} \quad (12)$$

(11) is obtained by replacing the class responsibilities  $\tau_{jk}^{(t)}$  by their unnormalized form, given in (9) in the summation. This operation leads to an expression that is exactly proportional to (10) when using hard assignments, but not for soft assignments. However, in the latter case, it generally has a very small effect on the final class responsibilities. Those are obtained by normalizing (12), as shown in Algorithm 3.

Although the three-step methods are generally dominated by the two-step approach regarding the bias and variance of the SM's parameters estimates (Bakk and Kuha 2018), they can present significant methodological advantages over the two-step approach for practitioners. Indeed, since the third step only depends on the MM through imputed class weights, estimating different SMs based on a three-step method does not require the researchers to have access to the indicator variables once the weights have been computed. For example, this allows the data owner to share only the relevant covariates, distal outcomes, and precomputed weight vectors with researchers interested in a particular structural model without disclosing the measurement model data. This property can be especially advantageous when the observed indicators used to estimate the measurement model contain sensitive data.

Regarding the differences between the naive and bias-adjusted three-step methods, the sole advantage of the naive three-step approach is that its interpretation is very intuitive due to its explicit classification step. The BCH method is very similar in this regard, except that it produces imputed class weights that can be negative. Although slightly more complex to interpret and computationally more expensive than the other three-step approaches, the ML method generally produces the best estimators of the SM's parameters (Vermunt 2010; Bakk *et al.* 2013; Bakk and Kuha 2018). Computational experiments carried out by Bakk *et al.* (2013) indicate that soft assignments generally produce better estimates than modal assignments for all the three-step methods, although the difference in the resulting SM's parameters estimates is usually relatively small.

**Algorithm 3** Three-step method

---

Initialize  $(\hat{\theta}_M^{(0)}, \hat{\theta}_S^{(0)}) \in (\Theta_M \times \Theta_S)$ , select  $\mathcal{A} \in \{\text{soft, modal}\}$ , select  $\mathcal{C} \in \{\text{ML, BCH, none}\}$

---

**STEP 1 : EM algorithm on the MM**

Set  $t = 0$

**repeat**

$$\begin{aligned} \tau_{ik}^{(t)} &= p(X_i=k|y_i; \hat{\theta}_M^{(t)}), \quad \forall i \in N_1, \forall k \in \mathcal{X} && \text{E Step} \\ \hat{\theta}_M^{(t+1)} &= \arg \max_{\theta_M \in \Theta_M} \sum_{i \in N_1} \omega_i \sum_{k \in \mathcal{X}} \tau_{ik}^{(t)} \log p(X_i=k, y_i; \theta_M) && \triangleright \text{Update responsibilities} \\ t &= t + 1 && \text{M Step} \\ & & & \hat{\theta}_M^{(t+1)} = \arg \max_{\theta_M \in \Theta_M} \sum_{i \in N_1} \omega_i \sum_{k \in \mathcal{X}} \tau_{ik}^{(t)} \log p(X_i=k, y_i; \theta_M) && \triangleright \text{Update } \hat{\theta}_M \end{aligned}$$

**until** convergence

Set  $\hat{\theta}_M = \hat{\theta}_M^{(t)}$   $\triangleright$  Estimated MM parameters

---

**STEP 2 : Class weights based on the MM**

Set  $w_{jk} = \begin{cases} p(X_j=k|y_j; \hat{\theta}_M) & , \text{ if } \mathcal{A} = \text{soft} \\ \mathbb{1}[\arg \max_{c \in \mathcal{X}} p(X_j=c|y_j; \hat{\theta}_M) = k] & , \text{ if } \mathcal{A} = \text{modal} \end{cases}, \forall j \in N_2, \forall k \in \mathcal{X}$   $\triangleright$

Weights

**if**  $\mathcal{C} \in \{\text{BCH, ML}\}$  **then**  $\triangleright$  Bias correction methods

    Define the r.v.  $W_j$  with pmf  $p(W_j=k) = w_{jk}, \forall j \in N_2, \forall k \in \mathcal{X}$   $\triangleright$  Predicted class membership

    Set  $D_{ck} = p(W=k|X=c; \hat{\theta}_M), \forall (c, k) \in \mathcal{X}^2$   $\triangleright$  Estimated misclassification probabilities

**if**  $\mathcal{C} = \text{BCH}$  **then**

$w_j = w_j D^{-1}, \forall j \in N_2$   $\triangleright$  Bias-adjusted weights

**else**

        Set  $w_j^* = w_j D^T, \forall j \in N_2$   $\triangleright$  Class-conditional probability of predicted class membership

**end if**

**end if**

---

**STEP 3 : EM algorithm on the SM, based on imputed class weights**

Set  $t = 0$

**repeat**

$$\begin{aligned} & & \text{E Step} \\ \text{if } \mathcal{C} = \text{ML} \text{ then} & & \\ \tau_{jk}^{(t)} &= w_{jk}^* p(X_j=k|z_j^o; z_j^p, \hat{\theta}_S^{(t)}), \quad \forall j \in N_2, \forall k \in \mathcal{X} && \triangleright \text{Update responsibilities} \\ \text{else} & & \\ \tau_{jk}^{(t)} &= w_{jk}, \quad \forall j \in N_2, \forall k \in \mathcal{X} && \triangleright \text{Imputed class weights used as responsibilities} \\ \text{end if} & & \\ \tau_{jk}^{(t)} &= \frac{\tau_{jk}^{(t)}}{\sum_{c \in \mathcal{X}} \tau_{jc}^{(t)}} \quad \forall j \in N_2, \forall k \in \mathcal{X} && \triangleright \text{Normalize responsibilities} \end{aligned}$$

$$\begin{aligned} & & \text{M Step} \\ \hat{\theta}_S^{(t+1)} &= \arg \max_{\theta_S \in \Theta_S} \sum_{j \in N_2} \omega_j \sum_{k \in \mathcal{X}} \tau_{jk}^{(t)} \log p(X_j=k, z_j^o; z_j^p, \theta_S) && \triangleright \text{Update } \hat{\theta}_S \\ t &= t + 1 && \end{aligned}$$

**until** convergence

Set  $\hat{\theta}_S = \hat{\theta}_S^{(t)}$   $\triangleright$  Estimated SM parameters

---

Return  $(\hat{\theta}_M, \hat{\theta}_S)$   $\triangleright$  Return estimated CM parameters

---

## 5. The StepMix package

This section gives an overview of the **StepMix** package. Before detailing the supported estimators and providing code examples, we first discuss the **StepMix** documentation and how our package relies on the existing Python ecosystem.

### 5.1. Software and documentation

**StepMix** follows the object-oriented interface of the **scikit-learn** Python library (Pedregosa *et al.* 2011), therefore exposing an API that is both familiar to machine learning practitioners and intuitive for beginners and students. The strict adherence to the **scikit-learn** interface enables easy comparisons with other latent class/clustering methods in the **scikit-learn** ecosystem and grants **StepMix** users access to cross-validation iterators for model selection and hyperparameter tuning.

The **StepMix** source code was heavily inspired by, and depends on, the **scikit-learn** Gaussian mixture implementation. Most computations rely on vectorized operations over **NumPy** arrays (Harris *et al.* 2020), and **StepMix** estimators can be saved and loaded using standard Python tools, such as the **Pickle** module. Documentation follows the **NumPy**-style docstring convention and is compiled into a web page at <https://StepMix.readthedocs.io/en/latest/>, which also features interactive notebook tutorials accessible to any user with a web browser.

### 5.2. Estimators

The simplest **StepMix** models do not include external variables (covariates and distal outcomes) and only have a single set of observed indicators. These models correspond to Gaussian or categorical mixtures, the latter of which is not currently available in **scikit-learn** (as of version 1.2). For example, assuming some integer-encoded categorical data is stored in the array  $Y$ , either as a **NumPy** array or a **pandas** data frame (McKinney 2010; The **pandas** Development Team 2020), a **StepMix** estimator for three latent classes can be declared and fit with the following commands.

```
model = StepMix(n_components = 3, measurement = "categorical")
model.fit(Y)
```

The most important parameters of a **StepMix** estimator are the number of latent classes (**n\_components**) and the specified conditional distribution (**measurement**). All options and distributions for the **measurement** argument are detailed in Table 2. Additional parameters of interest include seeding (**random\_state**) and various EM optimization parameters, such as the maximum number of EM iterations (**max\_iter**), the tolerance for stopping the optimization (**abs\_tol**) and the number of different initializations to try (**n\_init**). Users can obtain a detailed model report presenting parameters and fit statistics (**verbose=1**), examples of which are provided in Appendix A.

Adding a second set of observed variables and specifying a SM is a simple matter of adding a **structural** argument (the options of which are also detailed in Table 2). For the SM, **StepMix** relies on the familiar **scikit-learn** supervised learning interface: The SM data is provided as a second set of variables, typically reserved for target labels. This allows a clear

StepMix	Observed variable	Parameters (for each $x \in \{1, \dots, K\}$ )	PDF/PMF
binary	$z^o \in \{0, 1\}^D$	$\pi_x \in [0, 1]^D$	$p(z^o x) = \prod_{d=1}^D z_d^o \pi_{x,d} + (1 - z_d^o)(1 - \pi_{x,d})$
categorical	$z^o \in \{0, \dots, C\}^D$	$\mathcal{P}_x \in [0, 1]^{D \times C}$ $\sum_{c=1}^C \mathcal{P}_{x,d,c} = 1$	$p(z^o x) = \prod_{d=1}^D \mathcal{P}_{x,d,z_d^o}$
gaussian_unit	$z^o \in \mathbb{R}^D$	$\mu_x \in \mathbb{R}^D$	$p(z^o x) = \mathcal{N}(z^o; \mu_x, I)$
gaussian_spherical	$z^o \in \mathbb{R}^D$	$\mu_x \in \mathbb{R}^D, \sigma_x^2 \in \mathbb{R}$	$p(z^o x) = \mathcal{N}(z^o; \mu_x, \sigma_x^2 I)$
gaussian_diag	$z^o \in \mathbb{R}^D$	$\mu_x \in \mathbb{R}^D, \sigma_x^2 \in \mathbb{R}^D$	$p(z^o x) = \mathcal{N}(z^o; \mu_x, \text{diag}(\sigma_x^2))$
gaussian_full	$z^o \in \mathbb{R}^D$	$\mu_x \in \mathbb{R}^D, \Sigma_x \in \mathbb{R}^{D \times D}$	$p(z^o x) = \mathcal{N}(z^o; \mu_x, \Sigma_x)$
covariate	$z^p \in \mathbb{R}^D$	$\beta_x \in \mathbb{R}^D, b_x \in \mathbb{R}$	$p(x z^p) = \frac{\exp(\beta_x^\top z^p + b_x)}{\sum_{k=1}^K \exp(\beta_k^\top z^p + b_k)}$

Table 2: Probability distributions supported by **StepMix**, and the associated strings to be passed as the `measurement` and/or `structural` arguments of the constructor (first column) for some observed variable of dimension  $D$  and  $K$  latent classes. **StepMix** can be used to fit a single set of observed variables (by only specifying `measurement`) and provides a syntax for combining continuous and categorical variables in a single model. For notational clarity, parameters for a single latent class are shown, and all  $D$  categorical variables are integer-encoded.  $\mathcal{N}(\cdot; \mu, \Sigma)$  denotes the pdf of a Gaussian distribution parametrized by a mean  $\mu$  and a covariance matrix  $\Sigma$ .

separation between the MM and the SM data, and makes a clear distinction in the case where only a MM is specified.

The specific stepwise estimation procedure (`n_steps`) (see Section 4) and three-step specific arguments (`assignment`, `correction`) can also be passed. A **StepMix** model for three-step estimation with a BCH correction and soft assignment of a model with categorical indicators `Y` and continuous distal outcomes `Z_o` would look like:

```
model = StepMix(n_components = 3, measurement = "categorical",
                 structural = "gaussian_unit", n_steps = 3, assignment = "soft",
                 correction = "BCH")
model.fit(Y, Z_o)
```

As with any **scikit-learn** estimator, **StepMix** objects come with an array of useful methods for inference and model interpretation:

- `model.predict(Y, Z_o)` to predict the latent class  $x$  for each observation;
- `model.predict_proba(Y, Z_o)` to predict  $p(x | y, z_o)$  for each observation;
- `model.score(Y, Z_o)` to compute the average log-likelihood of the CM over the dataset;

- `model.aic(Y, Z_o)` to compute the Akaike information criterion (Akaike 1974) over the dataset;
- `model.bic(Y, Z_o)` to compute the Bayesian information criterion (Schwarz 1978) over the dataset;
- `model.get_mm_df()` to get the MM model parameters  $\theta_Y$  as a data frame;
- `model.get_sm_df()` to get the SM model parameters  $\theta_S$  as a data frame;
- `model.get_cw_df()` to get the marginal distribution over latent classes  $p(x; \theta_X)$  as a data frame;
- `model.sample(100)` to sample 100 observations from the fitted model.

The methods can be similarly used with some covariate data `Z_p` or some complete model data `Z`. We note that sampling from a fitted **StepMix** model is only possible if the marginal over the latent variable  $X$  is explicitly specified, i.e., when no covariates are used in the model (see (2)).

### 5.3. Nonparametric bootstrapping

Fitted **StepMix** objects include a `bootstrap_stats` method to obtain mean and standard error estimates of model parameters via nonparametric bootstrapping. We run a permutation test to align classes with the main estimator and avoid label switching between repetitions. The method can be called with

```
stats_dict = model.bootstrap_stats(Y, Z_o, n_repetitions = 10,
                                    progress_bar = True)
```

which returns a Python dictionary with the following attributes

- `stats_dict["samples"]`: Bootstrapped samples in a long-form format;
- `stats_dict["mm_mean"]`: Means of the MM parameters;
- `stats_dict["mm_std"]`: Standard errors of the MM parameters;
- `stats_dict["sm_mean"]`: Means of the SM parameters;
- `stats_dict["sm_std"]`: Standard errors of the SM parameters;
- `stats_dict["cw_mean"]`: Means of the class weights (marginal distribution);
- `stats_dict["cw_std"]`: Standard errors of the class weights.

All attributes in the `bootstrap_stats` output dictionary are **pandas** data frames.

Users can pass the data frame of bootstrapped samples to the **seaborn** library (Waskom 2021) for visualization, as shown in one of our tutorials (<https://colab.research.google.com/drive/14DJCqFTUaYp3JtLAeAMYmGHFLCHE-r7z>).

## 5.4. Additional features

**StepMix** provides additional functionalities, which are detailed in the documentation and tutorials, including:

- Support for **missing values**. Most models introduced in Table 2 can be suffixed with `_nan` (e.g., `binary_nan`) to enable full information maximum likelihood training (Section 4.1). The `gaussian_full` and `covariate` models do not currently support FIML;
- A dictionary-based syntax for combining multiple models in `measurement` or `structural`, allowing, for example, to declare a MM with **both categorical and continuous** variables or a SM with a covariate and an outcome (see example in Section 6.3).

Finally, **StepMix** was designed with modular methods for stepwise estimation, meaning each step in the estimation procedures can be represented as a standalone method call. For example, the naive three-step method can also be used by independently calling methods for i) fitting the MM with the EM algorithm, ii) performing soft assignments, and iii) fitting the SM by calling the M-step:

```
model.em(Y)
soft_assignments = model.predict_proba(Y)
model.m_step_structural(soft_assignments, Z_o)
```

Access to lower-level steps in the estimation will hopefully further stepwise estimation research by allowing researchers to investigate existing methods and more easily implement new ones.

## 6. Computational examples

In this section, we carry out two simulation studies to replicate **Latent Gold** results presented in [Bakk and Kuha \(2018\)](#) (Sections 6.1 and 6.1) followed by another simulation to test a complete model in the presence of missing values (Section 6.3). We then present a case study on a real-life dataset (Section 6.4) and conclude with a comparison of **StepMix** log-likelihoods with other packages on various datasets (Section 6.5).

The **StepMix** source code includes functions to simulate the datasets used in this section as well as instructions to reproduce our results. See the `scripts` directory at <https://github.com/Labo-Lacourse/stepmix>.

### 6.1. Distal outcome simulation study

For the first simulation study, we consider a dataset with  $K = 3$  latent classes,  $D_M = 6$  binary indicator variables, and  $D_S = 1$  continuous outcome variable, which is modeled as a mixture of unit variance Gaussians (see Table 2). Therefore, the MM's parameters are given by  $\theta_M = \{\rho, \pi\}$ , where  $\pi \in [0, 1]^{D_M \times K}$  represents class-conditioned binary probabilities, and the marginal distribution of the latent variable respects  $p(X=k) = \rho_k$ . The SM's parameters

are  $\theta_S = \{\mu\}$ ,  $\mu \in \mathbb{R}^{D_S \times K}$ . The true values of the parameters are:

$$\left\{ \begin{array}{l} \rho = [1/3, 1/3, 1/3], \quad \mu = [-1, 1, 0], \\ \\ \pi = \begin{bmatrix} \gamma & \gamma & 1-\gamma \\ \gamma & \gamma & 1-\gamma \\ \gamma & \gamma & 1-\gamma \\ \gamma & 1-\gamma & 1-\gamma \\ \gamma & 1-\gamma & 1-\gamma \\ \gamma & 1-\gamma & 1-\gamma \end{bmatrix}, \end{array} \right.$$

where the parameter  $\gamma$  controls the strength of the association between the latent class and the binary variables. [Bakk and Kuha \(2018\)](#) use for  $\gamma$  the values 0.9, 0.8, and 0.7 and call them the high-, medium-, and low-separation conditions, respectively.

To respect the original experiment design, we consider simulations with sample sizes  $n$  of 500, 1000, and 2000, resulting in nine configurations (three sample sizes  $\times$  three class separations). For each configuration, 500 data sets are generated. For simplicity, the authors only reported results for the estimates of parameter  $\mu_2 = 1$ . We therefore present the same statistics for all estimation methods. It should be noted that we use modal assignments for all three-step simulations following the specifications in [Bakk and Kuha \(2018\)](#). To minimize the variance in the estimation of the bias and the root-mean-squared error (RMSE) of the different estimators of  $\mu_2$  in different class separation settings, we used common random numbers to generate our simulated datasets.

The **StepMix** code for a single simulation of the high-separation one-step case is the following. Please note that the distal outcome is referred to as the `response` in the source code due to the terminology used in [Bakk and Kuha \(2018\)](#). We begin by simulating the dataset –  $Y$  is the matrix of MM data and  $Z_o$  is the matrix of SM outcome data – before instantiating the model and fitting it. We finally retrieve the means of the SM.

```
from stepmix.datasets import data_bakk_response
from stepmix.stepmix import StepMix

Y, Z_o, _ = data_bakk_response(n_samples = 2000, sep_level = 0.9,
                                random_state = 42)
model = StepMix(n_components = 3, measurement = "binary",
                structural = "gaussian_unit", n_steps = 1, random_state = 42,
                verbose = 1)
model.fit(Y, Z_o)
mus = model.get_sm_df()
```

The model report for the above example is presented in Appendix A.1.

The full results of our simulations are shown in Table 3. The **StepMix** results are consistent with the ones presented in Table 4, taken from [Bakk and Kuha \(2018\)](#) in which models were estimated with **Latent GOLD** version 5.1. The most noticeable difference is that our implementations of the two-step and ML methods seem to produce slightly better estimates

Class separation	Sample size	Mean bias						RMSE			
		1-step		2-step		3-step		1-step		3-step	
		Naive	BCH	ML	Naive	BCH	ML	Naive	BCH	ML	
Low ( $\gamma = 0.7$ )	500	0.10	-0.19	-0.64	-0.38	-0.14	0.6	0.3	0.6	0.8	0.1
	1000	0.10	-0.07	-0.61	-0.23	-0.05	0.1	0.2	0.3	0.6	0.2
	2000	-0.00	-0.02	-0.57	-0.11	-0.01	0.8	0.7	0.8	0.4	0.7
Medium ( $\gamma = 0.8$ )	500	0.00	-0.03	-0.31	-0.04	-0.02	0.1	0.5	0.3	0.8	0.5
	1000	-0.00	-0.02	-0.29	-0.01	-0.01	0.8	0.9	0.0	0.2	0.9
	2000	-0.00	-0.01	-0.28	0.00	-0.00	0.6	0.7	0.9	0.9	0.7
High ( $\gamma = 0.9$ )	500	-0.00	-0.01	-0.08	0.00	-0.00	0.9	0.9	0.2	0.9	0.9
	1000	-0.00	-0.00	-0.08	0.00	0.00	0.7	0.6	0.0	0.6	0.6
	2000	-0.00	-0.00	-0.08	0.00	-0.00	0.5	0.4	0.9	0.4	0.4

Table 3: Simulation results for **StepMix** point estimates of one parameter in an outcome SM (with true value  $\mu_2 = 1$ ). As described in the main text, three degrees of class separation in the MM (lower is harder) and three sample sizes are tested. Each configuration is simulated 500 times and the average parameter bias and RMSE are reported for one-, two- and three-step estimation, as well as the BCH and ML bias correction methods in the three-step case.

Class separation	Sample size	Mean bias						RMSE			
		1-step		2-step		3-step		1-step		3-step	
		Naive	BCH	ML	Naive	BCH	ML	Naive	BCH	ML	
Low ( $\gamma = 0.7$ )	500	0.1	-0.32	-0.66	-0.38	-0.36	0.9	0.4	0.0	0.0	0.7
	1000	0.1	-0.16	-0.60	-0.22	-0.20	0.2	0.0	0.3	0.7	0.3
	2000	0.0	-0.08	-0.58	-0.12	-0.11	0.8	0.1	0.9	0.7	0.3
Medium ( $\gamma = 0.8$ )	500	0.0	-0.03	-0.31	-0.03	-0.03	0.1	0.3	0.3	0.6	0.3
	1000	0.0	-0.01	-0.29	0.10	-0.01	0.8	0.9	0.0	0.2	0.9
	2000	0.0	0.10	-0.29	0.00	0.00	0.5	0.7	0.9	0.9	0.7
High ( $\gamma = 0.9$ )	500	0.0	0.00	-0.08	0.00	0.00	0.8	0.9	0.2	0.9	0.9
	1000	0.1	0.10	-0.07	0.10	0.10	0.6	0.6	0.0	0.7	0.6
	2000	0.0	0.00	-0.08	0.00	0.00	0.4	0.4	0.9	0.5	0.4

Table 4: Simulation results for point estimates of one structural parameter of an outcome SM (with true value  $\mu_2 = 1$ ) over 500 simulated data sets. Results are from [Bakk and Kuha \(2018\)](#) and produced using **Latent GOLD**. Please refer to Table 3 for simulation details.

compared to the commercial implementation in the low-separation condition, especially with small sample sizes.

The general conclusions presented by [Bakk and Kuha \(2018\)](#) can be drawn from both Tables 3 and 4. The one-step estimator is globally unbiased and systematically has the lowest RMSE. The naive three-step estimator is severely biased and has the highest RMSE. Its bias decreases with increasing class separation and is essentially unaffected by sample size. As the class separation level and the sample size increase, the bias-adjusted three-step methods progressively reduce this bias.

The two-step estimator is consistently better than the bias-adjusted three-step estimators.

Its smaller RMSE suggests that avoiding the extra step of three-step estimation may improve the quality of the resulting estimator. This conclusion appears to be quite intuitive since the second step of the three-step methods does not add any information to the model that was not already encapsulated in the MLE  $\hat{\theta}_M$  obtained in step one. In the medium- and high-separation conditions, the two-step estimator also performs essentially as well as the one-step estimator, suggesting that there is little loss of efficiency from moving from one-step estimation to an appropriate multi-step approach in this setting.

## 6.2. Covariate simulation study

The second simulation is concerned with a structural covariate model where the latent class  $X$  is a response to an observed one-dimensional predictor  $Z^p$ . Specifically, the SM is a multinomial logistic model (Table 2, covariate). For the simulation at hand,  $Z^p$  is a uniformly sampled integer between 1 and 5. The ground truth parameters of the structural model are:

$$\begin{cases} \beta = [0.00, -1.00, 1.00], \\ b = [0.00, 2.35, -3.66], \end{cases}$$

with  $\theta_S = (\beta, b)$ . The intercepts  $b$  were tuned to ensure approximately equal class sizes when averaged over  $Z^p$ . We consider the same binary indicators  $\pi$  as in the outcome simulation (Section 6.1), and the parameters of the MM are now defined as  $\theta_M = (\pi)$ , since we ignore the class weights  $\rho$  to follow the conditional likelihood perspective presented in (2) for the covariate case. The simulation specifications are otherwise identical to the ones presented in Section 6.1.

Similarly to Section 6.1, we present the code for a single covariate simulation. We begin by simulating the dataset.

```
from stepmix.datasets import data_bakk_covariate
from stepmix.stepmix import StepMix
Y, Z_p, _ = data_bakk_covariate(n_samples = 2000, sep_level = 0.9,
random_state = 42)
```

Contrary to other models, the covariate model does not have a closed-form MLE for the M-Step and requires a numerical optimization method for fitting parameters, the arguments of which can be optionally specified by passing a Python dictionary to the `structural_params` parameter in `StepMix`.

```
covariate_params = {"method": "newton-raphson", "max_iter": 1,
"intercept": True}
model = StepMix(n_components = 3, measurement = "binary",
structural = "covariate", n_steps = 1, random_state = 42,
structural_params = covariate_params)
model.fit(Y, Z_p)
```

We conclude by retrieving the coefficients of the SM and using the `pandas` library API to subtract the coefficients of the second class from all coefficients, therefore setting this class as the reference class.

Class separation	Sample size	Mean bias						RMSE					
		1-step		2-step		3-step		1-step		2-step		3-step	
		Naive	BCH	ML	Naive	BCH	ML	Naive	BCH	ML	Naive	BCH	ML
Low ( $\gamma = 0.7$ )	500	0.03	-0.27	-0.63	-0.29	-0.31	0.27	0.41	0.65	0.57	0.45		
	1000	0.01	-0.15	-0.61	-0.19	-0.18	0.17	0.29	0.62	0.42	0.33		
	2000	-0.00	-0.09	-0.62	-0.11	-0.10	0.12	0.20	0.62	0.33	0.23		
Medium ( $\gamma = 0.8$ )	500	0.02	-0.03	-0.39	-0.03	-0.05	0.16	0.18	0.40	0.29	0.21		
	1000	0.01	-0.02	-0.37	0.01	-0.02	0.12	0.13	0.38	0.21	0.14		
	2000	-0.00	-0.01	-0.37	0.01	-0.01	0.08	0.09	0.38	0.14	0.10		
High ( $\gamma = 0.9$ )	500	0.02	0.01	-0.12	0.01	0.01	0.13	0.13	0.16	0.15	0.14		
	1000	0.01	0.00	-0.12	0.01	0.01	0.09	0.09	0.15	0.11	0.10		
	2000	-0.00	-0.00	-0.13	-0.00	-0.00	0.07	0.07	0.14	0.07	0.07		

Table 5: Simulation results for **StepMix** point estimates of one parameter in a covariate SM (with true value  $\beta_2 = 1$ ) over 500 simulated data sets. Simulation details are presented in the main text and in Table 3.

Class separation	Sample size	Mean bias						RMSE					
		1-step		2-step		3-step		1-step		2-step		3-step	
		Naive	BCH	ML	Naive	BCH	ML	Naive	BCH	ML	Naive	BCH	ML
Low ( $\gamma = 0.7$ )	500	0.40	-0.24	-0.59	-0.25	-0.27	0.50	0.80	0.00	0.80	0.20		
	1000	0.30	-0.16	-0.60	-0.16	-0.20	0.90	0.20	0.20	0.50	0.50		
	2000	0.10	-0.09	-0.60	-0.11	-0.12	0.20	0.20	0.10	0.60	0.40		
Medium ( $\gamma = 0.8$ )	500	0.10	-0.05	-0.40	-0.03	-0.06	0.70	0.00	0.10	0.00	0.20		
	1000	0.10	-0.01	-0.37	0.20	-0.02	0.10	0.30	0.80	0.40	0.40		
	2000	-0.01	-0.01	-0.37	0.10	-0.01	0.80	0.90	0.80	0.50	0.00		
High ( $\gamma = 0.9$ )	500	0.20	0.10	-0.11	0.20	0.10	0.30	0.30	0.70	0.60	0.40		
	1000	0.10	0.10	-0.12	0.10	0.10	0.90	0.90	0.50	0.00	0.00		
	2000	0.00	0.00	-0.12	0.10	0.00	0.70	0.70	0.30	0.70	0.70		

Table 6: Simulation results for point estimates of one parameter in a covariate SM (with true value  $\beta_2 = 1$ ) over 500 simulated data sets. Results are from [Bakk and Kuha \(2018\)](#) and produced using **Latent GOLD**. Simulation details are presented in the main text and in Table 3.

```
betas = model.get_sm_df()
betas = betas.sub(betas[1], axis = 0)
```

Covariate simulation results are presented in Table 5. The results largely echo the ones presented in Section 6.1, with **StepMix** yielding by and large similar estimates to those found in [Bakk and Kuha \(2018\)](#) and reproduced in Table 6. Again, stepwise estimators (except naive three-step) produce globally unbiased estimates, and the **StepMix** implementation appears to slightly outperform **Latent GOLD** in the two-step and ML cases.

### 6.3. Complete simulation study

As a final simulation, we merge the outcome and covariate models of the previous sections

into a CM. We fix the degree of class separation in the MM ( $\gamma = 0.8$ ) and instead vary the ratio of data missing completely at random in both the MM and the outcome to showcase our FIML implementation.

We again begin our example by simulating the dataset. This time,  $Z$  includes the covariate variable as the first column and the outcome as the second column.

```
from stepmix.datasets import data_bakk_complete
from stepmix import StepMix
Y, Z, _ = data_bakk_complete(n_samples = 2000, sep_level = 0.8,
    nan_ratio = 0.25, random_state = 42)
```

While we previously specified models as simple strings (e.g., "binary"), the structural model for this example requires both a covariate and a continuous outcome. This setting requires a more complex model description, which with **StepMix** can be achieved using a nested Python dictionary. The nested dictionaries represent different submodels using different columns of  $Z$  as inputs. The keys in the main dictionary can be arbitrary strings to name the submodels.

The following descriptor indicates that the first column in  $Z$  should be passed to a **covariate** submodel with the given **method** and **max\_iter** arguments. The next column should be passed to a **gaussian\_unit\_nan** submodel. The **\_nan** suffix ensures support for missing values.

```
structural_descriptor = {
    "covariate": {"model": "covariate", "n_columns": 1,
        "method": "newton-raphson", "max_iter": 1},
    "response": {"model": "gaussian_unit_nan", "n_columns": 1}}
```

The above descriptor can be extended to accommodate an arbitrary number of columns and submodels. We could also define a similar **measurement\_descriptor**, although this is not needed here since the MM is homogenous and only includes **binary** variables. We conclude this example by instantiating a **StepMix** estimator and fitting it. Now that the structural model is more complex, we rely on the **pandas** API to specifically select the parameters of the SM model called **response**, as defined by the key in the descriptor.

```
model = StepMix(n_components = 3, measurement = "binary_nan",
    structural = structural_descriptor, n_steps = 1, random_state = 42)
model.fit(Y, Z)
mus = model.get_sm_df().loc["response"]
```

We present the complete simulation results in Table 7. We find that the one-step, two-step, and ML methods are globally unbiased, even in the presence of data missing completely at random. As expected, the RMSE increases with the ratio of missing values due to the decreasing effective sample size associated with each parameter. Naive three-step (and, to a lesser extent, three-step with BCH correction) yields noticeably inferior estimates as the ratio of missing data increases.

#### 6.4. Using parents' social status to predict respondents' family income

In this section, we illustrate the use of **StepMix** on a real-life dataset ( $n=3029$ ) that was previously studied by [Bakk, Oberski, and Vermunt \(2016\)](#) and [Bakk and Kuha \(2021\)](#). This

Missing values	Sample size	Mean bias						RMSE					
		1-step		2-step		3-step		1-step		2-step		3-step	
				Naive	BCH	ML				Naive	BCH	ML	
0%	500	0.01	-0.00	-0.30	-0.03	0.00	0.09	0.11	0.33	0.18	0.11		
	1000	-0.00	-0.00	-0.29	0.00	-0.00	0.06	0.07	0.30	0.13	0.08		
	2000	-0.00	-0.00	-0.28	0.01	-0.00	0.05	0.05	0.29	0.09	0.05		
25%	500	0.00	-0.04	-0.45	-0.13	-0.00	0.11	0.17	0.47	0.28	0.16		
	1000	-0.00	-0.01	-0.43	-0.08	-0.00	0.08	0.10	0.45	0.21	0.10		
	2000	0.00	0.00	-0.41	-0.03	0.00	0.05	0.07	0.42	0.14	0.07		
50%	500	0.00	-0.09	-0.57	-0.22	-0.00	0.14	0.24	0.60	0.40	0.20		
	1000	0.00	-0.04	-0.57	-0.16	0.00	0.10	0.17	0.59	0.31	0.15		
	2000	0.00	0.00	-0.54	-0.08	0.01	0.07	0.12	0.56	0.23	0.11		

Table 7: Simulation results for **StepMix** point estimates of a SM’s parameter in a model with six indicators, a covariate, and an outcome variable. We study the estimates of a single parameter in the outcome model (with true value  $\mu_2 = 1$ ) over 500 simulated data sets. For this simulation, we fix the degree of class separation in the MM ( $\gamma = 0.8$ ) and instead vary the ratio of data missing completely at random in the indicators and the outcome.

analysis is based on the 1976 ( $n_1=1499$ ) and 1977 ( $n_2=1530$ ) American General Social Surveys (GSS) conducted by the National Opinion Research Center. The data is publicly accessible via the **gssr** package (Healy 2025).

We consider a MM relating three indicators (the father’s job prestige, the mother’s highest degree and the father’s highest degree) to a latent variable with  $K = 3$  latent classes, interpreted as the parents’ social status (low, middle and high). The father’s job prestige score was measured on a scale ranging from 12 to 82 and recoded into three categories: Low (36 or less), medium (37-61) and high (62 or above). The father’s and mother’s education were measured on a 5-point scale ranging from 0 (less than high school) to 4 (graduate). In the SM, the parents’ social status is used to predict the respondent’s annual family income, in thousands of dollars. To account for the skewed distribution of this unique distal outcome, it is modeled as a normal variable with class-dependent mean and variance. Missing values for the annual family income are present in the GSS dataset.

This model was first estimated by Bakk *et al.* (2016) using the stepwise estimation methods implemented in **Latent GOLD** and handling the missing data through listwise deletion. Bakk and Kuha (2021) also compared the one-, two-, and three-step (BCH and ML) approaches on this dataset, but this time using FIML. We apply the same experimental setup as the latter study, but using **StepMix** instead of **Latent GOLD** for the estimation of each model. Appendix A.2 includes an example of a **StepMix** ML model report after fitting the GSS data.

Table 8 presents the estimated class proportions and the resulting class-conditional empirical distribution of the MM’s indicators. Table 9 reports the estimated parameters of the SM provided by each stepwise approach and their standard errors (SEs) and  $p$  values, estimated by nonparametric bootstrapping based on 100 repetitions. The reported SEs are obtained by calculating the standard deviation (SD) of a given bootstrapped parameter using StepMix’s nonparametric bootstrap module (Section 5.3). These bootstrapped parameters are the mean

Social status	Low	Middle	High
Class size	0.70	0.23	0.07
Father's job prestige			
Low	0.47	0.31	0.05
Medium	0.53	0.67	0.46
High	0.00	0.02	0.49
Mother's education			
Below high school	0.82	0.14	0.15
High school	0.17	0.79	0.44
Junior college	0.00	0.03	0.01
Bachelor	0.01	0.04	0.30
Graduate	0.00	0.01	0.10
Father's education			
Below high school	0.95	0.06	0.00
High school	0.05	0.89	0.11
Junior college	0.00	0.00	0.05
Bachelor	0.00	0.05	0.39
Graduate	0.00	0.00	0.44

Table 8: Estimated MM parameters (marginal distribution of social classes and class-conditional distribution of parents' education and job prestige).

Model	Low class income	Middle class income	High class income
1-step	20.32 (0.93)	26.10 (3.02)	67.90 (3.18)
2-step	25.25 (2.48)	38.41 (6.14)	50.66 (7.13)
Naive 3-step	27.44 (0.59)	35.94 (1.28)	43.67 (3.85)
3-step BCH	26.71 (0.74)	36.81 (1.64)	44.68 (4.26)
3-step ML	21.05 (1.52)	44.73 (5.25)	61.26 (11.77)

Table 9: Estimated SM parameters (class-conditional average annual family income) and their standard error for each stepwise method.

incomes for each latent class in Table 9 and the difference between the mean income of the reference latent class (low) and those of the middle and high classes in Table 10. Table 10 presents two-tailed  $Z$ -tests to assert whether the class-specific mean incomes for the middle and high classes are identical to that of the low social class.

The results obtained with **StepMix** are consistent with those reported by [Bakk and Kuha \(2021\)](#). Table 8 indicates that the low, middle and high social classes respectively compose approximately 70%, 23% and 7% of the sample. The class-conditional distribution of the indicators is highly heterogeneous, which highlights a strong association between the level of education and job prestige of the parents.

The results of Table 9 show that the estimated SM's parameters vary considerably between methods. For example, the estimated average family income in the highest social class ranges from 43.67 thousand dollars a year with the naive three-step method to 67.90 thousand dollars a year with the one-step method. In this case study, the assumption of normality of the class-conditional annual income is violated, as previously discussed by [Bakk \*et al.\* \(2016\)](#)

Model	Est.	SE	Z	$P(> z )$
1-step				
Middle class	5.79	3.78	1.53	$p = .126$
High class	47.58	3.49	13.63	$p < .001$
2-step				
Middle class	13.16	8.38	1.57	$p = .116$
High class	25.41	6.22	4.09	$p < .001$
Naive 3-step				
Middle class	8.50	1.41	6.02	$p < .001$
High class	16.22	3.83	4.24	$p < .001$
3-step BCH				
Middle class	10.11	1.93	5.24	$p < .001$
High class	17.97	4.22	4.26	$p < .001$
3-step ML				
Middle class	23.68	6.27	3.77	$p < .001$
High class	40.21	11.58	3.47	$p < .001$

Table 10: Family's income differences between classes for each method.

and [Bakk and Kuha \(2021\)](#). In the one-step approach, the joint estimation of the MM and the SM propagates this specification error from the SM to the definition of the latent classes. As a result, the estimated proportion of each class in the sample changes to 62%, 20% and 18%, and the empirical distribution of the indicators in both smaller classes is practically indistinguishable. This observation illustrates that the definition of the SM can significantly impact the interpretation of the latent class in the one-step approach, making it less suitable for real-life applications.

Regarding the variance of the estimators, the naive three-step method leads to the lowest bootstrap estimates of the SM's parameters' SEs. However, the resulting estimated class-conditional average incomes are biased since this method does not consider the uncertainty of class assignments. Among the two-step and bias-corrected three-step estimators, the BCH method has the smallest SEs across bootstrapped means. This is consistent with the results obtained by [Bakk and Kuha \(2021\)](#), whose results indicate that BCH method is the least sensitive to violations of distributional assumptions. Compared to the three-step ML approach, the two-step method has noticeably closer estimated means to the BCH method.

These differences in the estimated SM's parameters and their SEs are important from an applied perspective as they can lead to significant differences in the conclusion of the study. Based on the  $p$  values of the Z-tests of Table 10, researchers would reach different conclusions regarding the difference in respondent's family income between those whose parents were classified in the low and middle social statuses. According to the more robust BCH approach, we should conclude that the incomes are, on average, significantly higher when their parents had a middle social status ( $Z = 5.24; p < 0.001$ ) than low social status. Researchers using the one-step ( $Z = 1.53; p = 0.126$ ) or the two-step approach ( $Z = 1.57; p = 0.116$ ) would incorrectly fail to reject the null hypothesis. This highlights the importance of comparing different stepwise estimators when studying real data, as imperfect model specification is almost unavoidable in practice.

Software	Latent class analysis			Latent profile analysis		
	Carcinoma <i>MM only</i>	Response simulation <i>1-step</i>	Covariate simulation <i>1-step</i>	Iris <i>MM only</i>	Diabetes <i>1-step</i>	Banknote <i>1-step</i>
<b>StepMix</b>	−293.705	−8600.793	−5192.049	−180.185	−2407.146	−771.669
<b>poLCA</b>	−293.705	−	−5192.049	−	−	−
<b>mclust</b>	−	−	−	−180.186	−	−
<b>MoEClust</b>	−	−	−	−180.186	−	−771.669
<b>Mplus</b>	−293.705	−8599.609	−5192.049	−180.185	−2407.146	−771.668

Table 11: Comparison of **StepMix** log-likelihoods with log-likelihoods from **poLCA** (Linzer and Lewis 2011), **mclust** (Scrucca *et al.* 2016), **MoEClust** (Murphy and Murphy 2020), and **Mplus** (Muthén and Muthén 2017). For latent class analysis (binary or categorical MM), we consider the Carcinoma dataset (Agresti 2002) and the two simulation datasets from Bakk and Kuha (2018) discussed in Sections 6.1 and 6.2. For latent profile analysis (continuous MM), we benchmark the Iris (Anderson 1935; Fisher 1936), Diabetes (Reaven and Miller 1979), and Banknote (Flury 1988) datasets. We used three classes for each dataset except for Banknote, where we used two. Latent class models were computed with default MM parameters. In the case of latent profile models, we used diagonal covariance matrices for each class in the Diabetes and Banknote examples, which corresponds to **mclust**’s “VVI” model and **StepMix**’s **gaussian\_diag** model. For the Iris example, we used general covariance matrices for each class, which corresponds to **mclust**’s “VVV” model and **StepMix**’s **gaussian\_full** model. For **Mplus**, we manually defined the covariance structure to match other packages. The 1-step approach was used for all SMs to facilitate package comparison. Results were obtained using the default optimization parameters of each package, except in the Iris example, where 20 random initializations were used for **StepMix** and **Mplus**. The “class” variable was used as a categorical distal outcome in the Diabetes example, and “status” was used as a binary covariate in the Banknote example. All packages found solutions in under one second for all datasets. The “−” entries indicate models that are not supported and were thus not estimated by the corresponding package.

## 6.5. Comparison with other packages

We present a comparison of the log-likelihoods obtained with **StepMix** and four other mixture modeling packages on six datasets in Table 11. **StepMix** yields log-likelihoods that are nearly identical to existing commercial and open-source software, validating the quality of our implementation.

## 7. Conclusion

We presented **StepMix**, an open-source package for mixture model estimation with continuous and categorical variables. To the best of our knowledge, **StepMix** is the first Python package to implement bias-adjusted stepwise estimators for mixture models with external variables, potentially exposing a completely new user base to these methods. The familiar **scikit-learn** interface, the support for missing values, and the R wrapper are additional features that we believe will make **StepMix** a relevant tool for the community. Furthermore, we presented all estimation procedures as variants of the EM algorithm to provide a unified framework of

analysis. This framework also guided the implementation of **StepMix**, resulting in modular source code that is conducive to the implementation of new features, such as novel stepwise estimation methods or additional conditional distributions. As a future development, all of the 14 Gaussian covariance decompositions implemented in **mclust** and **MoEClust** could be added to **StepMix**. We could also implement efficient initialization methods for the EM algorithm instead of relying on random initialization. It would also be interesting to extend the capabilities of **StepMix** to deal with high dimensional data using recent regularization techniques (Robitzsch 2020), and potentially handle large data sets through the implementation of additional EM methods. Other extensions could include the support for generalized mixture models for longitudinal data, such as growth mixture models (Ram and Grimm 2009), latent class growth models (Nagin 1999), and mixtures of hidden Markov models (Van de Pol and Langeheine 1990).

## Acknowledgments

Sacha Morin and Robin Legault contributed equally to this work.

This research was funded in part by Natural Sciences and Engineering Research Council of Canada (NSERC) PGS D Scholarships [SM, RL], by Fonds de recherche du Québec – Nature et technologies (FRQNT) master’s [SM, RL] and doctoral scholarships [S.M., R.L.], by IVADO MSc Excellence Scholarships [SM, RL], as well as by the Canadian Institutes of Health Research (CIHR) (grant number 170633), the Social Sciences and Humanities Research Council (SSHRC) and the Centre for the Study of Democratic Citizenship (CSDC) [RS]. This research was enabled in part by computing resources provided by Mila (mila.quebec).

## References

Agresti A (2002). *Categorical Data Analysis*. John Wiley & Sons.

Akaike H (1974). “A New Look at the Statistical Model Identification.” *IEEE Transactions on Automatic Control*, **19**(6), 716–723. [doi:10.1109/tac.1974.1100705](https://doi.org/10.1109/tac.1974.1100705).

Anderson E (1935). “The Irises of the Gaspé Peninsula.” *Bulletin of American Iris Society*, **59**, 2–5. [doi:10.1130/spe117-p1](https://doi.org/10.1130/spe117-p1).

Asparouhov T, Muthén B (2014). “Auxiliary Variables in Mixture Modeling: Three-Step Approaches Using **Mplus**.” *Structural Equation Modeling: A Multidisciplinary Journal*, **21**(3), 329–341. [doi:10.1080/10705511.2014.915181](https://doi.org/10.1080/10705511.2014.915181).

Athey TL, Liu T, Pedigo BD, Vogelstein JT (2019). “**Autogmm**: Automatic and Hierarchical Gaussian Mixture Modeling in **Python**.” *arXiv 1909.02688*, arXiv.org E-Print Archive. [doi:10.48550/arXiv.1909.02688](https://doi.org/10.48550/arXiv.1909.02688).

Bakk Z, Kuha J (2018). “Two-Step Estimation of Models Between Latent Classes and External Variables.” *Psychometrika*, **83**(4), 871–892. [doi:10.1007/s11336-017-9592-7](https://doi.org/10.1007/s11336-017-9592-7).

Bakk Z, Kuha J (2021). “Relating Latent Class Membership to External Variables: An Overview.” *British Journal of Mathematical and Statistical Psychology*, **74**(2), 340–362. [doi:10.1111/bmsp.12227](https://doi.org/10.1111/bmsp.12227).

Bakk Z, Oberski DL, Vermunt JK (2016). “Relating Latent Class Membership to Continuous Distal Outcomes: Improving the Ltb Approach and a Modified Three-Step Implementation.” *Structural Equation Modeling*, **23**(2), 278–289. [doi:10.1080/10705511.2015.1049698](https://doi.org/10.1080/10705511.2015.1049698).

Bakk Z, Tekle FB, Vermunt JK (2013). “Estimating the Association between Latent Class Membership and External Variables Using Bias-Adjusted Three-Step Approaches.” *Sociological Methodology*, **43**(1), 272–311. [doi:10.1177/0081175012470644](https://doi.org/10.1177/0081175012470644).

Bandeen-Roche K, Miglioretti DL, Zeger SL, Rathouz PJ (1997). “Latent Variable Regression for Multiple Discrete Outcomes.” *Journal of the American Statistical Association*, **92**(440), 1375–1386. [doi:10.1080/01621459.1997.10473658](https://doi.org/10.1080/01621459.1997.10473658).

Bartolucci F, Montanari GE, Pandolfi S (2014). “A Comparison of Some Estimation Methods for Latent Markov Models with Covariates.” In *Proceedings of COMPSTAT*, pp. 531–538.

Beath KJ (2017). “**randomLCA**: An R Package for Latent Class with Random Effects Analysis.” *Journal of Statistical Software*, **81**(13), 1–25. [doi:10.18637/jss.v081.i13](https://doi.org/10.18637/jss.v081.i13).

Besag J (1975). “Statistical Analysis of Non-Lattice Data.” *Journal of the Royal Statistical Society D*, **24**(3), 179–195. [doi:10.2307/2987782](https://doi.org/10.2307/2987782).

Bolck A, Croon M, Hagenaars J (2004). “Estimating Latent Structure Models with Categorical Variables: One-Step versus Three-Step Estimators.” *Political Analysis*, **12**(1), 3–27. [doi:10.1093/pan/mp001](https://doi.org/10.1093/pan/mp001).

Croon M (2002). “Using Predicted Latent Scores in General Latent Structure Models.” In G Marcoulides, I Moustaki (eds.), *Latent Variable and Latent Structure Modeling*, pp. 195–223. Erlbaum.

Dayton CM, Macready GB (1988). “Concomitant-Variable Latent Class Models.” *Journal of the American Statistical Association*, **83**, 173–178. [doi:10.1080/01621459.1988.10478584](https://doi.org/10.1080/01621459.1988.10478584).

Dempster AP, Laird NM, Rubin DB (1977). “Maximum Likelihood From Incomplete Data via the EM Algorithm.” *Journal of the Royal Statistical Society B*, **39**(1), 1–22. [doi:10.1111/j.2517-6161.1977.tb01600.x](https://doi.org/10.1111/j.2517-6161.1977.tb01600.x).

Devlieger I, Mayer A, Rosseel Y (2016). “Hypothesis Testing Using Factor Score Regression: A Comparison of Four Methods.” *Educational and Psychological Measurement*, **76**(5), 741–770. [doi:10.1177/0013164415607618](https://doi.org/10.1177/0013164415607618).

Fisher RA (1936). “The Use of Multiple Measurements in Taxonomic Problems.” *Annals of Eugenics*, **7**(2), 179–188. [doi:10.1111/j.1469-1809.1936.tb02137.x](https://doi.org/10.1111/j.1469-1809.1936.tb02137.x).

Flury B (1988). *Multivariate Statistics: A Practical Approach*. Chapman and Hall.

Fraley C, Raftery AE (2007). “Bayesian Regularization for Normal Mixture Estimation and Model-Based Clustering.” *Journal of Classification*, **24**, 155–181. [doi:10.1007/s00357-007-0004-5](https://doi.org/10.1007/s00357-007-0004-5).

Geman S, Geman D (1984). “Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **6**, 721–741. [doi:10.7551/mitpress/4943.003.0038](https://doi.org/10.7551/mitpress/4943.003.0038).

Goldberger AS (1961). “Stepwise Least Squares: Residual Analysis and Specification Error.” *Journal of the American Statistical Association*, **56**(296), 998–1000. [doi:10.1080/01621459.1961.10482142](https://doi.org/10.1080/01621459.1961.10482142).

Gong G, Samaniego FJ (1981). “Pseudo Maximum Likelihood Estimation: Theory and Applications.” *The Annals of Statistics*, pp. 861–869. [doi:10.1214/aos/1176345526](https://doi.org/10.1214/aos/1176345526).

Grün B, Leisch F (2008). “**FlexMix** Version 2: Finite Mixtures with Concomitant Variables and Varying and Constant Parameters.” *Journal of Statistical Software*, **28**(4), 1–35. [doi:10.18637/jss.v028.i04](https://doi.org/10.18637/jss.v028.i04).

Harris CR, Millman KJ, Van Der Walt SJ, Gommers R, Virtanen P, Cournapeau D, Wieser E, Taylor J, Berg S, Smith NJ, *et al.* (2020). “Array Programming with **NumPy**.” *Nature*, **585**(7825), 357–362. [doi:10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2).

Hartley HO, Hocking RR (1971). “The Analysis of Incomplete Data.” *Biometrics*, pp. 783–823. [doi:10.2307/2528820](https://doi.org/10.2307/2528820).

Healy K (2025). **gssr**: General Social Survey Data for Use in R. R package version 0.7, URL <https://kjhealy.github.io/gssr>.

Kim Y, Jeon S, Chang C, Chung H (2022). “**glca**: An R Package for Multiple-Group Latent Class Analysis.” *Applied Psychological Measurement*, **46**(5), 439–441. [doi:10.1177/01466216221084197](https://doi.org/10.1177/01466216221084197).

Koller D, Friedman N (2009). *Probabilistic Graphical Models: Principles and Techniques*. MIT Press.

Lacourse E, Baillargeon R, Dupéré V, Vitaro F, Romano E, Tremblay R (2010). “Two-Year Predictive Validity of Conduct Disorder Subtypes in Early Adolescence: A Latent Class Analysis of a Canadian Longitudinal Sample.” *Journal of Child Psychology and Psychiatry*, **51**(12), 1386–1394. [doi:10.1111/j.1469-7610.2010.02291.x](https://doi.org/10.1111/j.1469-7610.2010.02291.x).

Lebret R, Iovleff S, Langrognet F, Biernacki C, Celeux G, Govaert G (2015). “**Rmix-mod**: The R Package of the Model-Based Unsupervised, Supervised, and Semi-Supervised Classification **Mixmod** Library.” *Journal of Statistical Software*, **67**(6), 1–29. [doi:10.18637/jss.v067.i06](https://doi.org/10.18637/jss.v067.i06).

Linzer DA, Lewis JB (2011). “**poLCA**: An R Package for Polytomous Variable Latent Class Analysis.” *Journal of Statistical Software*, **42**(10), 1–29. [doi:10.18637/jss.v042.i10](https://doi.org/10.18637/jss.v042.i10).

Lyrvall J, Mari RD, Bakk Z, Oser J, Kuha J (2023). “**multilevLCA**: An R Package for Single-Level and Multilevel Latent Class Analysis with Covariates.” *arXiv 2305.07276*, arXiv.org E-Print Archive. [doi:10.48550/arXiv.2305.07276](https://doi.org/10.48550/arXiv.2305.07276).

Marbac M, Sedki M (2019). “**VarSelLCM**: An R/C++ Package for Variable Selection in Model-Based Clustering of Mixed-Data with Missing Values.” *Bioinformatics*, **35**(7), 1255–1257. [doi:10.1093/bioinformatics/bty786](https://doi.org/10.1093/bioinformatics/bty786).

McKinney W (2010). “Data Structures for Statistical Computing in Python.” In S Van der Walt, J Millman (eds.), *Proceedings of the 9th Python in Science Conference*, pp. 56–61.

McLachlan GJ, Krishnan T (2007). *The EM Algorithm and Extensions*. John Wiley & Sons.

McLachlan GJ, Lee SX, Rathnayake SI (2019). “Finite Mixture Models.” *Annual Review of Statistics and Its Application*, **6**(1), 355–378. [doi:10.1146/annurev-statistics-031017-100325](https://doi.org/10.1146/annurev-statistics-031017-100325).

Meyer D, Dimitriadou E, Hornik K, Weingessel A, Leisch F (2024). **e1071**: *Misc Functions of the Department of Statistics, Probability Theory Group (Formerly: E1071)*, TU Wien. [doi:10.32614/CRAN.package.e1071](https://doi.org/10.32614/CRAN.package.e1071). R package version 1.7-16.

Mulder E, Vermunt J, Brand E, Bullens R, Van Merle H (2012). “Recidivism in Subgroups of Serious Juvenile Offenders: Different Profiles, Different Risks?” *Criminal Behaviour and Mental Health*, **22**, 122?135. [doi:10.1002/cbm.1819](https://doi.org/10.1002/cbm.1819).

Murphy K, Murphy TB (2020). “Gaussian Parsimonious Clustering Models with Covariates and a Noise Component.” *Advances in Data Analysis and Classification*, **14**(2), 293–325. [doi:10.1007/s11634-019-00373-8](https://doi.org/10.1007/s11634-019-00373-8).

Muthén LK, Muthén B (2017). *Mplus User’s Guide: Statistical Analysis with Latent Variables*.

Nagin DS (1999). “Analyzing Developmental Trajectories: A Semiparametric, Group-Based Approach.” *Psychological Methods*, **4**(2), 139. [doi:10.1037/1082-989X.4.2.139](https://doi.org/10.1037/1082-989X.4.2.139).

Oberski D (2016). “Mixture Models: Latent Profile and Latent Class Analysis.” In *Modern Statistical Methods for HCI*, pp. 275–287. Springer-Verlag, Cham. [doi:10.1007/978-3-319-26633-6\\_12](https://doi.org/10.1007/978-3-319-26633-6_12).

Okazaki S, Campo S, Andreu L, Romero J (2015). “A Latent Class Analysis of Spanish Travelers’ Mobile Internet Usage in Travel Planning and Execution.” *Cornell Hospitality Quarterly*, **56**(2), 191–201. [doi:10.1177/1938965514540206](https://doi.org/10.1177/1938965514540206).

Ormerod JT, Wand MP (2010). “Explaining Variational Approximations.” *The American Statistician*, **64**(2), 140–153. [doi:10.1198/tast.2010.09058](https://doi.org/10.1198/tast.2010.09058).

Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay E (2011). “**scikit-learn**: Machine Learning in Python.” *Journal of Machine Learning Research*, **12**, 2825–2830. URL <https://www.jmlr.org/papers/v12/pedregosa11a.html>.

Ram N, Grimm KJ (2009). “Methods and Measures: Growth Mixture Modeling: A Method for Identifying Differences in Longitudinal Change among Unobserved Groups.” *International Journal of Behavioral Development*, **33**(6), 565–576. [doi:10.1177/0165025409343765](https://doi.org/10.1177/0165025409343765).

R Core Team (2025). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. [doi:10.32614/R.manuals](https://doi.org/10.32614/R.manuals). URL <https://www.R-project.org/>.

Reaven GM, Miller RG (1979). “An Attempt to Define the Nature of Chemical Diabetes Using a Multidimensional Analysis.” *Diabetologia*, **16**, 17–24. [doi:10.1007/bf00423145](https://doi.org/10.1007/bf00423145).

Robitzsch A (2020). “Regularized Latent Class Analysis for Polytomous Item Responses: An Application to SPM-LS Data.” *Journal of Intelligence*, **8**(3), 30. [doi:10.3390/jintelligence8030030](https://doi.org/10.3390/jintelligence8030030).

Schwarz G (1978). “Estimating the Dimension of a Model.” *The Annals of Statistics*, pp. 461–464. [doi:10.1214/aos/1176344136](https://doi.org/10.1214/aos/1176344136).

Scrucca L, Fop M, Murphy TB, Raftery AE (2016). “**mclust** 5: Clustering, Classification and Density Estimation Using Gaussian Finite Mixture Models.” *The R Journal*, **8**(1), 289. [doi:10.32614/rj-2016-021](https://doi.org/10.32614/rj-2016-021).

Sterba SK (2013). “Understanding Linkages among Mixture Models.” *Multivariate Behavioral Research*, **48**(6), 775–815. [doi:10.1080/00273171.2013.827564](https://doi.org/10.1080/00273171.2013.827564).

Stroustrup B (2013). *The C++ Programming Language*. 4th edition. Addison-Wesley.

The **pandas** Development Team (2020). “**pandas-dev/pandas: pandas**.” [doi:10.5281/zenodo.3509134](https://doi.org/10.5281/zenodo.3509134).

Van de Pol F, Langeheine R (1990). “Mixed Markov Latent Class Models.” *Sociological Methodology*, pp. 213–247. [doi:10.2307/271087](https://doi.org/10.2307/271087).

Van Rossum G, Drake FL (2009). *Python 3 Reference Manual*. CreateSpace, Scotts Valley.

Vermunt JK (2010). “Latent Class Modeling with Covariates: Two Improved Three-Step Approaches.” *Political Analysis*, **18**(4), 450–469. [doi:10.1093/pan/mpq025](https://doi.org/10.1093/pan/mpq025).

Vermunt JK, Magidson J (2013). *Technical Guide for Latent GOLD 5.0: Basic, Advanced, and Syntax*. Statistical Innovations, Belmont.

Visser I, Speekenbrink M (2010). “**depmixS4**: An R Package for Hidden Markov Models.” *Journal of Statistical Software*, **36**(7), 1–21. [doi:10.18637/jss.v036.i07](https://doi.org/10.18637/jss.v036.i07).

Waskom ML (2021). “**seaborn**: Statistical Data Visualization.” *Journal of Open Source Software*, **6**(60), 3021. [doi:10.21105/joss.03021](https://doi.org/10.21105/joss.03021).

White A, Murphy TB (2014). “**BayesLCA**: An R Package for Bayesian Latent Class Analysis.” *Journal of Statistical Software*, **61**(13), 1–28. [doi:10.18637/jss.v061.i13](https://doi.org/10.18637/jss.v061.i13).

Wu CFJ (1983). “On the Convergence Properties of the EM Algorithm.” *The Annals of Statistics*, **11**(1), 95–103. ISSN 00905364. [doi:10.1214/aos/1176346060](https://doi.org/10.1214/aos/1176346060).

Xue QL, Bandeen-Roche K (2002). “Combining Complete Multivariate Outcomes with Incomplete Covariate Information: A Latent Class Approach.” *Biometrics*, **58**(1), 110–120. [doi:10.1111/j.0006-341x.2002.00110.x](https://doi.org/10.1111/j.0006-341x.2002.00110.x).

## A. StepMix output

### A.1. Distal outcome simulation study

By setting `verbose=1`, **StepMix** outputs a full model report, including parameter estimates and fit statistics. The output for the code example in Section 6.1 is as follows.

```

Fitting StepMix...
=====
MODEL REPORT
=====
=====

Measurement model parameters
=====

  model_name      binary
  class_no          0      1      2
  param variable
  pis   feature_0  0.9028  0.0928  0.9024
        feature_1  0.8996  0.1044  0.9108
        feature_2  0.8978  0.1024  0.9039
        feature_3  0.1156  0.0832  0.8992
        feature_4  0.0756  0.0885  0.9034
        feature_5  0.1004  0.0962  0.9080
=====

Structural model parameters
=====

  model_name      gaussian_unit
  class_no          0      1      2
  param variable
  means feature_0      0.9599 -0.0379 -1.0153
=====

Class weights
=====

  Class 1 : 0.34
  Class 2 : 0.34
  Class 3 : 0.32
=====

Fit for 3 latent classes
=====

  Estimation method      : 1-step
  Number of observations : 2000
  Number of latent classes : 3
  Number of estimated parameters: 23
  Log-likelihood (LL)      : -8600.7930
  -2LL                      : 17201.5860
  Average LL                : -4.3004
  AIC                      : 17247.59
  
```

```

BIC : 17376.41
CAIC : 17399.41
Sample-Size Adjusted BIC : 17478.16
Entropy : 183.3425
Scaled Relative Entropy : 0.9166

```

## A.2. Application example

StepMix output for the ML estimator on the GSS data presented in Section 6.4.

```

Fitting StepMix...
=====
MODEL REPORT
=====
=====

Measurement model parameters
=====

      model_name          categorical_nan
      class_no
      param variable
      pis   Father's education_0      0.9452  0.0553  0.0000
            Father's education_1      0.0533  0.8863  0.1137
            Father's education_2      0.0015  0.0050  0.0530
            Father's education_3      0.0000  0.0534  0.3899
            Father's education_4      0.0000  0.0000  0.4434
            Father's job prestige_0    0.4667  0.3056  0.0497
            Father's job prestige_1    0.5304  0.6735  0.4616
            Father's job prestige_2    0.0029  0.0209  0.4887
            Mother's education_0      0.8215  0.1376  0.1523
            Mother's education_1      0.1658  0.7857  0.4352
            Mother's education_2      0.0039  0.0295  0.0107
            Mother's education_3      0.0076  0.0399  0.3012
            Mother's education_4      0.0012  0.0073  0.1006
=====

Structural model parameters
=====

      model_name          gaussian_diag_nan
      class_no
      param variable
      covariances Income (1000) 144.7759  607.0978 2633.7187
      means      Income (1000)  20.7196   45.2885  66.0730
=====

Class weights
=====

  Class 1 : 0.70
  Class 2 : 0.23

```

```

Class 3 : 0.07
=====
Fit for 3 latent classes
=====
Estimation method      : 3-step
Correction method      : ML
Assignment method      : modal
Number of observations  : 3029
Number of latent classes: 3
Number of estimated parameters: 38
Log-likelihood (LL)      : -18754.7237
-2LL                      : 37509.4475
Average LL                : -6.1917
AIC                      : 37585.45
BIC                      : 37814.05
CAIC                     : 37852.05
Sample-Size Adjusted BIC : 37997.92
Entropy                   : 624.9492
Scaled Relative Entropy   : 0.8122

```

## B. R code

### B.1. Distal outcome simulation

This section replicates the Python example from section 6.1 using the R interface **stepmixr**.

```

R> library("stepmixr")
R> datasim <- data_bakk_response(n_samples = 2000, sep_level = 0.9,
+      random_state = 42)
R> Y = datasim[[1]]
R> Z_o = datasim[[2]]
R> model = stepmix(n_components = 3, measurement = "binary",
+      structural = "gaussian_unit", n_steps = 1, random_state = 42,
+      verbose = 1)
R> fit1 = fit(model, Y, Z_o)
R> mus = fit1$get_sm_df()

```

### B.2. Covariate simulation

This section replicates the Python example from section 6.2 using the R interface **stepmixr**.

```

R> library("stepmixr")
R> datasim <- data_bakk_covariate(n_samples = 2000, sep_level = 0.9,
+      random_state = 42)

```

```
R> Y <- datasim[[1]]
R> Z_p <- datasim[[2]]
R> covariate_params <- list(method = "newton-raphson", max_iter = as.integer(1),
+    intercept = TRUE)
R> model <- stepmix(n_components = 3, measurement = "binary",
+    structural = "covariate", n_steps = 1, random_state = 42,
+    structural_params = covariate_params)
R> fit1 <- fit(model, Y, Z_p)
R> betas <- fit1$get_parameters()
R> betas <- betas$structural$beta
R> betas <- betas - c(1, 1, 1) %o% betas[2, ]
```

### B.3. Complete simulation study

This section replicates the Python example from section 6.3 using the R interface **stepmixr**.

```
R> library("stepmixr")
R> datasim = data_bakk_complete(n_samples = 2000, sep_level = 0.8,
+    nan_ratio = 0.25, random_state = 42)
R> Y <- datasim[[1]]
R> Z <- datasim[[2]]
R> structural_descriptor <- list(
+    covariate = list(model = "covariate", n_columns = as.integer(1),
+        method = "newton-raphson", max_iter = as.integer(1)),
+    response = list(model = "gaussian_unit_nan", n_columns = as.integer(1)))
R> model <- stepmix(n_components = 3, measurement = "binary_nan",
+    structural = structural_descriptor, n_steps = 1, random_state = 42)
R> fit1 <- fit(model, Y, Z)
R> mus <- fit1$get_parameters()
R> mus <- mus$structural$response
```

### Affiliation:

Éric Lacourse  
 Département de sociologie  
 Faculté des arts et des sciences  
 Université de Montréal  
 C.P. 6128, succursale Centre-ville  
 Montréal, Québec, Canada  
 H3C 3J7  
 E-mail: [eric.lacourse@umontreal.ca](mailto:eric.lacourse@umontreal.ca)