Latent Class Growth Analysis and Growth Mixture Modeling using R: A tutorial for two R-packages and a comparison with Mplus.

Working paper (Work in progress: updated and corrected version 5;
December 2022)

Klaas J. Wardenaar

Department of Psychiatry,
University Medical Center Groningen
Groningen, the Netherlands

Summary

Latent Class Growth Analyses (LCGA) and Growth Mixture Modeling (GMM) analyses are used to explain between-subject heterogeneity in growth on an outcome by identifying latent classes with different growth trajectories. Dedicated software packages are available to estimate these models, with *Mplus* (Muthén & Muthén, 2019) being widely used¹. Although this and other available commercial software packages are of good quality, very flexible and rich in options, they can be costly and fit poorly into the analytical workflow of researchers that increasingly depend on the open-source R-platform. Interestingly, although plenty of R-packages to conduct mixture analyses are available, there is little documentation on how to conduct LCGA/GMM in R. This educational document aims to provide applied researchers with relatively little prior experience with a tutorial and coding examples for conducting LCGA and GMM in R. The results obtained with R and the modeling approaches (e.g., default settings, model configuration) of the used R-packages are compared to each other and to Mplus using simulated data.

Corrections and additions in this version (v5, Dec 2022)

In this revised version of this document I added a foreword as well as additional clarifications, recommendations and references throughout the document. Furthermore, I corrected the layout of Table 1, Table 3, and spelling/grammar/typos throughout the document. More specifically, I changed the description of how to correctly use the proportionality coefficient when using the 'lcm' package. Based on this, some variance-covariance parameters in Table 3 have been corrected as well. Finally, I included sections to stress the importance of multiple random starts when running LCGA/GMM with multiple starts and added some clarifications with regard to missing data handling of the different discussed packages.

Correspondence

K.J. Wardenaar PhD
University Medical Center Groningen
Department of Psychiatry (CC-72)
Hanzelplein 1
9713 GZ Groningen
Email: k.j.wardenaar@umcg.nl

¹ Some major statistical software packages also offer the possibility to estimate LCGA and GMM: SAS (Proc Traj), Stata/GLLAMM (Stata corp, TX),

Foreword (Dec 2022)

I initially wrote this document as an educational handout for my students and for researchers in my department. Since uploading it, I have found that others found it helpful as well. Readers should be advised that this educational document remains a work in progress. It is not (planned to be) published in a peer-reviewed journal and should not be treated nor cited as such.

This work is aimed for readers with a average to good understanding of statistics, but with little prior experience with more complex modeling and R. The work should be read as an introduction to get started and not as a comprehensive work on the theory and practice of conducting LCGA and/or GMM. After getting started, most researchers will be required to take additional steps and conduct more extensive modeling to answer their research question. A wealth of information on LCGA/GMM and how to properly use it for research can be found in the literature. Some useful references are given below, but much more is available.

Finally, make sure to always consult the documentation that comes with the software/package version that you use. As software is updated over time, the way in which models are fitted, default settings, and available options may change. In addition, keep in mind that not all available functions of each software/package are covered in this document, whereas some of the functions that are not covered here may be highly relevant to your specific research. The versions of Mplus and the R-packages used to prepare this document have mostly been superseded by newer versions. The functionality of the given code-examples has not been checked for these newer versions and cannot be guaranteed.

All that being said, I hope this document will be helpful to get started with LCGA/GMM in R.

Introduction

Latent Class Growth Analyses (LCGA) and Growth Mixture Modeling (GMM) are widely used statistical models in social, behavioural and medical science. They can be used to identify latent subgroups, classes or clusters of individuals based on their common growth trajectories over time. This tutorial focuses on illustrating the procedures to carry out the actual analyses, and thus assumes basic knowledge on the topics of longitudinal data analysis and linear mixed models (i.e. the distinction between fixed and random effects). Absolute beginners could start by reading up on growth curve modeling (e.g., Curran et al., 2010) and LCGA/GMM (e.g., Jung & Wickrama, 2008; Ram & Grimm, 2009; van de Schoot, 2015; van de Schoot et al., 2017; van der Nest et al., 2020) before going further and starting any actual analyses.

It is useful to look briefly at what LCGA and GMM are and how they are related. Both techniques are related and it can be helpful to regard both as extended versions of *growth models*. Growth models are used to quantify temporal trends or growth patterns in longitudinal data. Often, such models describe the outcome as a linear function of a time-variable, but differently shaped functions are also possible (e.g., quadratic, logarithmic, cubic). The most restrictive type of growth model focuses on estimating the average growth, quantified by a fixed intercept and slope describing the average trend over time in the data. The resulting fixed-effects model has the advantage of being easy to interpret and estimate, but often shows suboptimal fit to the data, especially in cases where substantial between-subject heterogeneity in change over time exists. It is possible to address this limitation by using less restrictive growth models that account for between-subject heterogeneity by including random effects (e.g., a random intercept and, sometimes, random slope). This means that on top of the fixed effects (intercepts and slope), normally distributed variances around the fixed intercept and/or slope are estimated. This effectively means that a distribution of individual growth trajectories instead of a single mean trend are estimated, accounting for between-subject heterogeneity in growth and allowing for explanation of between-subject variance (e.g., by including covariates). Such random effects models² allow for much more flexibility, explain both within- and between-subject variance in the outcome, and often show better statistical fit to real-life data than fixed effect models. However, some disadvantages of random effects models are that they are more complex (more parameters) and harder to estimate than fixed-effect models.

For educational purposes, it is helpful to look at LCGA as an extension of a fixed-effects growth model and GMM as an extension of a random-effects growth model. Both techniques belong to the *latent class* or *mixture* model family, where each latent class has distinct growth characteristics (i.e. growth model parameters). As it allows for within-class variation of the growth-model parameters (i.e. intercepts and slopes), GMM is more flexible than LCGA. The latter assumes that no within-class variance and covariance exist (this is called the assumption of local independence) and estimates fixed growth-model parameters in each class. As a result of these differences in flexibility, GMM models often require less latent classes than LCGA to explain the same heterogeneity in observed growth trajectories. This will also become apparent when working through the examples in this document.

It is important to note that the terms ‘latent class’ and ‘mixture’ both refer to the concept that the total observed data consists of a finite number of mixtures or latent (‘invisible’) classes, each of which is characterized by more homogeneity in the parameters or of the underlying model³. LCGA is essentially a very restrictive version of a GMM. Mixture models are used in many other statistical applications as well, ranging from relatively simple (e.g., Latent Class Analysis (LCA; Lazarsfeld, 1950; Lazarsfeld & Henry, 1968) to complex (e.g., Mixed Measurement Item Response Theory (MM-IRT; Mislevy & Verhelst, 1990)). Irrespective of the exact type of used mixture model, the procedure that is followed to identify the optimal number of classes is roughly the same. Models with increasing numbers of classes are estimated and model-fit is compared between the subsequent models to select the model that best describes the data.

This document focuses on the use of the R-platform (R-core Team, 2010) for estimating LCGA and GMM. R has several advantages over commercially available software, such as Mplus. R-packages are free and all code is open-source, making the software broadly accessible and transparent. In addition, the R-platform and (many) packages are continually maintained/improved by a very active community of researchers and programmers. A downside is that the learning-curve for R is steep and a huge number of packages exist (of varying quality). The latter makes it hard to know where to start and what package to use for a given task, especially for beginning users. In addition, for less experienced analysts it can be a challenge to find out if a chosen package actually does what they want it to do. R-packages generally tend to cater to more advanced-level analysts: procedures like LCGA/GMM can be carried out, but are often not as thoroughly documented as for purpose-built software like Mplus. R-documentation in general can be quite minimal or technical.

This educational document aims to provide applied researchers with information and coding examples needed to get started with using R to estimate LCGA/GMM models. The use of two user friendly R-packages that both offer mixture growth modeling functionality as part of their extensive set of options is demonstrated. In addition, as software packages often differ in terms of the configuration or operationalization of the estimated model(s) and with regard to default settings, these factors are compared between the packages and their effects on the results evaluated. Finally, the general consistency of obtained results across R-packages and Mplus is evaluated.

² Confusingly, many different terms are used in the literature to refer to such models, including Multilevel Models, Hierarchical Linear Models and Mixed Effects Models.

³ For novice users it may be very useful to know from the start that in statistical jargon, the term *mixture model* characterizes models used to identify latent classes (by parametric methods), whereas the term *mixed model* refers to models that include both fixed and random effects. Confusingly, GMM with random effects can in fact be seen as a type of 'Mixture Mixed Model'.

1. Methods

1.1 The dataset

The used dataset (download [here](#)) dataset contains simulated data for 100 cases, each with 5 equally spaced repeated measures on a continuous outcome scale (y1-y5). The simulations are based on a linear mixed model growth model with a random intercept and random slope. The data are simulated with a grouping variable that splits the sample into two (quite obvious) classes of 50 subjects with differing mean growth trajectories. The simulation R-code is provided in **Appendix 1**.

Figure 1 clearly shows that in the simulated data there is considerable heterogeneity in growth, both across and between classes. Note that these data are model-driven and merely used for illustration of the software functionality. They are not representative of typically observed growth-trajectory data. Real data are usually messier, with more outliers, non-linear trends, unequal follow-up time-intervals, and missing values. Also, unlike in the simulated example data, there usually is no visible distinction between classes in observed data. Even without the colors, two general growth-trajectory classes can be distinguished in **Figure 1**. In practice, such obvious separation is seldom encountered. This makes identification of the optimal model based on real observed data much harder and leaves more room for ambiguity (e.g., with regard to the optimal number of classes to capture the most important growth variations). Of course, unlike in our simulated case, class-membership is never known a priori.

With regard to structure, data can have a ‘long’ format, with multiple rows per subject and the number of rows per subjects corresponding to the number of repeated measurements, or a ‘wide’ format, with a single row per subject and a variable (column) for each repeated measurement of the outcome. The R-packages use the long format and Mplus uses the wide format.

Figure 1: simulated data for use in the tutorial: the different colors denote the two simulated subgroups.

1.2 Software

There are several commercially available software packages to conduct LCGA or GMM. Probably the most widely used packages are SAS (PROC TRAJ (Jones et al., 2001; mainly for LCGA) and Mplus (Muthén & Muthén, 2013; for both LCGA and GMM). Mplus offers a broad range of modeling options, covering most structural equation modeling approaches, latent variable models and various item response theory models. Given that Mplus is so widely used for LCGA and GMM, it will be used to provide the comparison models in the current paper⁴.

For R, there are several packages that allow for growth modeling and/or mixture modeling. Of

these, several allow users to fit LCGA and/or GMM models, of which two packages will be illustrated here in more detail. First, the ‘*lcmm*’ package offers many advanced options for mixture modeling with longitudinal data, but also provides the options needed to conduct LCGA and GMM, using the ‘*hlme*’ function (Proust-Lima et al., 2017; <https://cran.r-project.org/web/packages/lcmm/index.html>). Second, the ‘*flexmix*’ package (Leisch, 2004; <https://cran.r-project.org/web/packages/flexmix/index.html>) provides functionality for flexible mixture regression modeling, including LCGA and GMM models. The models are run in R using R-studio. The used versions for this tutorial are Mplus 5.0, R v3.6.2, R-studio v1.2.5033, ‘*lcmm*’ v1.8.1, and ‘*flexmix*’ v2.3-15.

When working on your own analyses, make sure to check which version you are using and always consult the latest version of the package’s reference manual/documentation. Updates and bug-fixes may eventually necessitate alterations to the code-examples provided below.

1.3 Covered models and analytical approach

This tutorial shows how to get started with the following model variants: (1) LCGA (fixed intercept and slope), (2) Growth Mixture Modeling with class-specific random intercepts (GMM-1), and (3) Growth Mixture Modeling with class-specific random intercepts and random slopes (GMM-2). The estimated models will be limited to linear growth functions.

Because there is always a risk of a mixture-model solution at a local rather than a global maximum in mixture modeling, all models are run multiple times (100 times) with varying random start values. Across these runs, the solution with the highest log likelihood is identified. This is important. Make sure to consult relevant literature to make sure you understand the importance of using a sufficient number of multiple random starts. Hipp and Bauer (2006) provide an excellent in depth discussion. In the provided examples, 100 random starts are used for illustration, but in practice you may want to use more random starts, depending – among other things - on the complexity of the fitted models.

With each package, the same general analytical approach is used. In each analysis, the first step of the process is to define the growth model. Second, models with increasing numbers of classes (1-3 classes) are fit to the data. Third, the optimal model is identified by comparing the Bayesian Information Criterion (BIC) and Akaike Information Criterion (AIC) across the fitted models, with the lowest values indicating the best model given the data. Finally, the number of subjects allocated to each class based on their highest posterior class probabilities are inspected and the estimated growth-trajectory parameters are inspected.

Given that the data are simulated based on a 2-class model with a random intercept and a random slope, a 2-class GMM-2 model is expected to show the best fit in the examples. Note that in real world applications the data-generation process is unknown and the optimal model type can be different. The optimal model configuration (LCGA, GMM-1, GMM-2 or other variant) will eventually depend on the data and needs to be established as part of the data-driven analysis process.

1.4 Code examples

For each package (*lcmm*, *flexmix* and *Mplus*), example code is provided in dedicated text boxes. Users need to make sure to adjust data file names, object names, and variable names according to their specific setting. For clarity, all user-specified elements (datasets, objects and variable names) are printed in blue. All annotation throughout the code is printed in green. The examples do only cover the basic function-arguments. Consult the package documentation for all available and possibly relevant options and settings.

A final note of caution: there is no simple one-size-fits-all optimal approach to conducting most statistical procedures and this also does apply here. The provided code can be used to get started, but researchers need to make sure that they (a) understand what the code actually does and (b) need to decide on necessary subsequent steps themselves. Here, the optimal next steps (e.g., increasing the number of random starts, accounting for non-linear growth, inclusion of covariates etc.) will depend on their specific research question and dataset.

⁴ It is important to point out that there are other commercial packages also offer possibilities to estimate LCGA, GMM and more. These include Latent Gold (Vermunt & Magidson, 2016) and Stata (GLLAMM; Rabe-Hesketh et al., 2004).

Table 1: model fit statistics for latent class growth models (LCGA) estimated with different software packages using the same simulated dataset.

Software	Number of classes	Free parameters	LL	AIC	BIC	Class-sizes
'lcm'	1-class	3	-1146.3	2298.6	2306.5	100
	2-class	6	-875.5	1762.9	1778.5	50, 50
	3-class	9	-858.3	1734.6	1758.0	48, 19, 33
'flexmix' ²	1-class	3	-1146.3	2298.6	2311.3	100
	2-class	6	-875.5	1762.9	1788.2	50, 50
	3-class	9	-858.3	1734.6	1772.5	19, 48, 33
Mplus ¹	1-class	3	-1146.3	2298.6	2306.5	100
	2-class	6	-875.5	1762.9	1778.5	50, 50
	3-class	9	-858.3	1734.6	1758.0	33, 48, 19

All models fitted with 100 random starts.

¹This model was run while overriding the default setting to estimate a residual variance term per time point by imposing an equality constraint across time-points.

²BIC calculated based on total number of data points instead of total number of subjects.

Table 2: model fit statistics for growth mixture models (GMM), estimated with different software packages using the same simulated dataset.

Software		Number of classes	Free parameters	Log-likelihood	AIC	BIC	Class-sizes
'lcm'	Random intercept	1-class	4	-1042.0	2091.9	2102.4	100
		2-class	8	-855.3	1726.7	1747.5	50, 50
		3-class	12	-838.7	1701.4	1732.7	35, 45, 20
	Random intercept + slope	1-class	6	-857.5	1727.1	1742.7	100
		2-class	10	-828.8	1677.5	1703.6	50, 50
		3-class	14	-826.0	1679.9	1716.4	18, 50, 32
'flexmix' ²	Random intercept	1-class	4	-1042.0	2091.9	2108.8	100
		2-class	8	-855.3	1726.7	1760.4	50, 50
		3-class	12	-838.7	1701.4	1752.0	20, 35, 45
	Random intercept + slope	1-class	6	-857.5	1727.1	1752.4	100
		2-class	12	-828.4	1680.7	1731.3	50, 50
		3-class	18	-824.6	1685.1	1761.0	22, 49, 29
Mplus ¹	Random intercept	1-class	4	-1042.0	2091.9	2102.4	100
		2-class	8	-855.3	1726.7	1747.5	50, 50
		3-class	12	-838.7	1701.4	1732.7	35, 45, 20
	Random Intercept + slope	1-class	6	-857.5	1727.1	1742.7	100
		2-class	12	-829.4	1680.7	1712.0	50, 50
		3-class ³	18	-820.6	1677.1	1724.0	8, 42, 50

All models fitted with 100 random starts.

The lines in bold font represent the models with the best fit according to the AIC and/or BIC.

The models with a random intercept + slope take a long time to estimate due to model complexity and (in case of the 3-class model) due to model misspecification given the data.

¹This model was run while overriding the default setting to estimate a residual variance term per time point by imposing an equality constraint across time-points.

²BIC calculated based on total number of data points instead of total number of subjects.

³Mplus gives a warning that there are estimation difficulties (latent variable covariance matrix is not positive definite in the very small first class).

2. Tutorial and Results

2.1 Package *lcmm*

2.1.1 LCGA with *lcmm*

The `hlme()` function from the *lcmm* package can be used to run a LCGA in R, using the following code:

```
# set working directory (location of data)
set.wd(...)

# open package
library(lcmm)

# set the seed for random number generator, so results can be
# reproduced:
set.seed(2002)

# run models with 1-4 classes, each with 100 random starts,
# using the 1-class model to set initial start values:
lcga1 <- hlme(y ~ time, subject = "ID", ng = 1, data = mydata)
lcga2 <- gridsearch(rep = 100, maxiter = 10, minit = lcga1,
  m=hlme(y ~ time, subject = "ID",
    ng = 2, data = mydata, mixture = ~ time))
lcga3 <- gridsearch(rep = 100, maxiter = 10, minit = lcga1,
  m=hlme(y ~ time, subject = "ID",
    ng = 3, data = mydata, mixture = ~ time))

# make table with results for the 3 models:
summarytable(lcga1, lcga2, lcga3)
```

We start by opening the package and setting the seed that will be used by the random number generator. This allows us to reproduce the results of the procedures with multiple random starts if we rerun the analyses. The `hlme()` function starts with a formula statement (`y~time`) to define the growth model. The `subject="ID"` statement identifies the unit (subject) within which the repeated measures are nested. The function `hlme()` allows for specification of random effects (slope and intercept). In the code above, these are omitted, leading to estimation of LCGA models with a fixed intercept and slope per class. Here, `hlme()` is used as a sub-function in the `gridsearch()` function. The latter is used to run `hlme()` 100 times (`rep=100`), each time using different start values to avoid ending up on a local optimum. Here and below, 100 random starts are used for illustration. It is important to choose a sufficiently large number here (Hipp & Bauer, 2006) and you can (and should) increase this number, especially when fitting more complex models with more classes. The `maxiter` argument is used to set the maximum number of iterations in the optimization algorithm in each random start. Here it is set at 10, but it can also be increased if needed. In this case, the start values are based on the 1-class model (`minit=lcga1`). All of the user-specified settings for the LCGA and GMM models can of course be further adjusted to fit one's aims (see the package documentation for guidance). Running the code in the simulated data yields individual model objects (`lcga1-lcga3`). The `summarytable()` function can be used to tabulate all models for comparison on number of estimated parameters (`npm`), indices of fit (the best log likelihood value across random starts, AIC and BIC) and class sizes:

G		loglik	npm	AIC	BIC	%class1	%class2	%class3
lcga1	1	-1146.3214	3	2298.6432	306.458	100		
lcga2	2	-875.4555	6	1762.9111	778.542	50	50	
lcga3	3	-858.2967	9	1734.5931	758.040	48	19	33

To inspect a specific model more closely, we can look at its model summary. Here, we can for example type `'summary(lcga2)'` in the console, which provides the following, more extensive, information:


```

Goodness-of-fit statistics:
  maximum log-likelihood: -875.46
  AIC: 1762.91
  BIC: 1778.54

Maximum Likelihood Estimates:

Fixed effects in the class-membership model:
(the class of reference is the last class)

      coef      Se    wald p-value
intercept class1 -0.00075 0.12108 -0.006 0.99509

Fixed effects in the longitudinal model:

      coef      Se    waldp-value
intercept class1 10.49814 0.13299 78.9390.00000
intercept class2  9.95581 0.13297 74.8720.00000
time class1      1.76654 0.05439 32.4780.00000
time class2      0.26565 0.05432  4.8900.00000

coef                      Se Residual standard
error:                    1.21375 0.03844

```

In addition to the estimation details and goodness-of-fit information, this output also shows the class-specific intercepts and slopes (i.e. ‘time class1’ and ‘time class2’) as well as the residual standard error.

If we stopped here and only ran these LCGA models, we could be tempted to conclude that a 3-class LCGA model best describes our data, as this model has the lowest AIC and BIC values. However, we would have to estimate LCGA models with more classes (i.e. 4 and more) to make sure that a more complex model does not fit the data even better. Here it is likely that an LCGA model with a considerable number of classes is needed, given the heterogeneity of the growth trajectories in the data. Such a model would not be an accurate nor a parsimonious representation of the actual data (we know that the data contains two classes). Here, it is wise not to stop, but to move on to more flexible GMM.

2.1.2 GMM-1 with *lcmm*

To estimate GMMs with only a random intercept using *lcmm*, the LCGA code can be adjusted as follows:

```

set.seed(2002)
gmm1 <- hlme(y ~ time, subject = "ID", random=~1, ng = 1, data =
  mydata)
gmm2 <- gridsearch(rep = 100, maxiter = 10, minit = gmm1,
  hlme(y ~ time, subject = "ID", random=~1,
  ng = 2, data = mydata, mixture = ~ time,
  nwg=T))
gmm3 <- gridsearch(rep = 100, maxiter = 10, minit = gmm1,
  hlme(y ~ time, subject = "ID", random=~1,
  ng = 3, data = mydata, mixture = ~ time,
  nwg=T))

# make table with results for the 4 models:
summarytable(gmm1, gmm2, gmm3)

```

The main difference with the LCGA code lies in the addition of the ‘random=~1’ part to the model definition. This indicates that we want to include a random intercept per class: i.e. a fixed mean intercept with normally distributed intercept variance around it. The ‘nwg=T’ statement indicates that we want to allow the variances of the random intercepts to vary across classes. Again, the 1-class model is used for initial start values, each model is estimated 100 times using the ‘gridsearch()’ function and `maxiter` is set at 10. The `rep` and `maxiter` settings can be changed according to the user’s needs. The summary of GMM-1 results (`summarytable(gmm1, gmm2, gmm3)`) looks like this:

G		loglik	npm	AIC	BIC	%class1	%class2	%class3
gmm1	1	-1041.9742	4	2091.948	2102.369	100		
gmm2	2	-855.3405	8	1726.681	1747.522	50	50	
gmm3	3	-838.7191	12	1701.438	1732.700	35	45	20

Here, we can see that compared to the LCGA models, additional free parameters are estimated per model (one additional parameter per class). These are the class-specific random intercept variances/proportionality coefficients. We can also see that the AIC and BIC values are considerably lower than for the LCGAs, indicating that the GMMs describe the data considerably better, but we also see that the 2-class model is not the best GMM-1 model. A specific model can be inspected more closely, for instance, by typing `'summary(gmm2)'` in the console. This gives the following output:

```
Goodness-of-fit statistics:
  maximum log-likelihood: -855.34
  AIC: 1726.68
  BIC: 1747.52

Maximum Likelihood Estimates:

Fixed effects in the class-membership model:
(the class of reference is the last class)

      coef      se    wald p-value
intercept class1  0.00813 0.20634  0.0390.96856

Fixed effects in the longitudinal model:

      coef      se    wald    p-value
intercept class1  9.95919 0.14039 70.939 0.00000
intercept class2 10.49674 0.14642 71.687 0.00000
time class1      0.26815 0.04801  5.585 0.00000
time class2      1.76959 0.04843 36.539 0.00000

Variance-covariance matrix of the random-effects:
      intercept
intercept 0.3874

      coef      se
Proportional coefficient class1 0.88937 0.22389
Residual standard error:      1.06407 0.03767
```

The output is very similar to the LCGA output, with the only difference being the added estimation of a class-specific random intercept variance, displayed in the **'Variance-covariance matrix of the random-effects'**. In *lcmm*, the random effect variance is estimated for the highest numbered class (here: class 2) and the other class-specific variances are calculated by multiplying the estimated variance-covariance matrix by a proportional coefficient, which is provided at the bottom of the output (i.e. **'Proportional coefficient class1'**). How this works in practice is explained here by the package developers: <https://github.com/CecileProust-Lima/lcmm/issues/64>. We multiply the intercept variance for class 2 (0.3874) by the squared proportionality coefficient to obtain the intercept variance for class 1 ($0.8894^2 \times 0.3874$), which gives us roughly 0.31.

Similar to the LCGA models, the AIC and BIC values did not reach a minimum value for the 2-class GMM-1 model. The main reason for this, is that the data were simulated with random slopes, which is not accounted for in the GMM-1 models. With a GMM-1 model, more classes seem to be needed to capture the heterogeneity in the subjects' growth trajectories. A GMM-2 model is better suited in this context.

2.1.3 GMM-2 with lcmm

To run a GMM-2 with both random intercepts and slopes, the GMM-1 code can be adjusted as follows:

```
set.seed(2002)
gmm1_2 <- hlme(y ~ time, subject = "ID", random=~1 + time, ng = 1,
  data=mydata)
gmm2_2 <- gridsearch(rep = 100, maxiter = 10, minit = gmm1_2,
  hlme(y ~ time, subject = "ID", random=~1 + time,
    ng = 2, data = mydata, mixture = ~ time, nwg=T))
gmm3_2 <- gridsearch(rep = 100, maxiter = 10, minit = gmm1_2,
  hlme(y ~ time, subject = "ID", random=~1+time,
    ng = 3, data = mydata, mixture = ~ time,
    nwg=T))

# make table with results for the 4 models:
summarytable(gmm1_2, gmm2_2, gmm3_2)
```

Here, the term ‘random=~1+time’ indicates that we want to estimate a random intercept and slope (‘random=~time’ would accomplish the same; the intercept is included by default). All other arguments are kept the same (e.g., `rep`, `maxiter`), but can of course be adjusted according to the user’s needs. After running the model with our simulated data, the summarized results (‘`summarytable(gmm1_2, gmm2_2, gmm3_2)`’) look like this:

	G	loglik	npm	AIC	BIC	%class1	%class2	%class3
gmm1_2	1	-857.5434	6	1727.087	1742.718	100		
gmm2_2	2	-828.7719	10	1677.544	1703.595	50	50	
gmm3_2	3	-825.9719	14	1679.944	1716.416	18	50	32

We can see that each model has two more freely estimated parameters than the GMM with random intercept because a random slope and the covariance between the random intercept and random slope are estimated as well in this model configuration. Again, it can be seen that the model fit of these models is better than that of the less complex versions (LCGA and GMM-1). In addition, we can see that the 2-class model – unsurprisingly given the used population model for simulation - best describes the data, as indicated by the lowest AIC and BIC values (also see **Table 2**). In addition, the class-sizes remain the same compared to LCGA and GMM-1.

The model output for this 2-class model can be further inspected by typing ‘`summary(gmm2_2)`’ in the console. The output looks as follows:

```

Goodness-of-fit statistics:
  maximum log-likelihood: -828.77
    AIC: 1677.54
    BIC: 1703.6

Maximum Likelihood Estimates:

Fixed effects in the class-membership model:
(the class of reference is the last class)

      coef      se    wald p-value
intercept class1 0.00521 0.20572  0.025 0.97980

Fixed effects in the longitudinal model:

      coef      se    wald p-value
intercept class1 10.49948 0.12268 85.586 0.00000
intercept class2  9.95285 0.11975 83.113 0.00000
time class1      1.76169 0.06725 26.195 0.00000
time class2      0.26605 0.06194  4.295 0.00002

Variance-covariance matrix of the random-effects:
      intercept      time
intercept  0.18818
time      -0.05754 0.10088

      coef      se
Proportional coefficient class1 1.11805 0.22247
Residual standard error:      0.92713 0.03775

```

Here, the ‘Variance-covariance matrix of the random-effects’ has three entries: the random intercept variance, random slope variance and random intercept-slope covariance for the highest numbered class (class 2). Again, the corresponding values for class 1 are obtained by multiplying the class-2 matrix by the provided squared proportional coefficient (See <https://github.com/CecileProust-Lima/lcmm/issues/64> for an explanation of how this works).

The use of a proportional coefficient to calculate all remaining class-specific (co)variance matrices leads to GMM models with less parameters (only a proportional coefficient per added class) than in an arrangement where the variance-covariance matrix has to be estimated for each class separately (which is the approach of *Mplus* and *flexmix*). This becomes clear when comparing the numbers of free parameters of the GMM-1 and GMM-2 models between *lcmm* and the other packages (see **Table 2**). As a result of this approach, the estimated models in *lcmm* have less parameters, making computations more efficient and reducing the time needed to estimate the models. However, these models do make the assumption that all elements in each class-specific random effects matrix can indeed be adequately approximated by multiplication of the estimated matrix with a proportional coefficient. See McNeish et al. (2022) for a thorough discussion of these types of models.

2.2 Package *flexmix*

2.2.1 LCGA in *flexmix*

To run LCGA models with increasing numbers of classes in the '*flexmix*' package, we can run the '*LXMRglmfix()*' function within the '*stepFlexmix()*' function. The latter allows for automatically stepping through LCGA models with increasing numbers of classes. The code looks as follows:

```
# open package:
library(flexmix)

# set random number seed:
set.seed(0111)

# estimate lcg models with 1-3 classes:
lcgaMix <- stepFlexmix(~ .|ID, k = 1:3, nrep = 100, model
= FLXMRglmfix(y~ time, varFix=TRUE), data = mydata, control
= list(iter.max = 500, minprior = 0))

# inspect and compare model-fit results:
lcgaMix

# extract a specific model object (2-class model) from the
# lcgaMix object:
lcga2 <- getModel(lcgaMix, which=2)

# inspect the overall model and parameters of the 2-class
# model:
summary(lcga2)
parameters(lcga2)

# calculate class sizes by dividing the provided sizes by
# the number of time points per subject, for example:
lcgaMix@models$`2`@size/5
```

We start by opening the package (*library(flexmix)*). As with the previous examples, we set the seed for the random number generator (*set.seed(...)*). The *stepFlexmix()* function is used to allow for estimation of models with multiple random starts. We start by defining the model formula. Here, a default *~.* is entered, which is updated with the formula provided in the *model* statement. The added element '*| ID*' indicates that measures are nested within subjects. The *k=1:3* statement indicates that we want to estimate models with 1 to 3 classes, and '*nrep=100*' indicates that we want to run each model 100 times with random starts to prevent local maxima. In this and following examples, 100 random starts are used for illustration. As noted above, it is important to choose a sufficiently large number here (Hipp & Bauer, 2006) and the number should be increased, especially when fitting more complex models. The *model* statement is used to specify the growth model. Here, we make use of the function *FLXMRglmfix()*, which allows users to specify general linear models for mixture analyses, while fixing certain parameters to be equal across classes. Here, we specify that we want the residual term to be equal across classes by *varFix=TRUE*. This leads to the estimation of a LCGA with a single residual term that reflects the overall unexplained variance in the data and does not vary across classes. This is analogous to the default in *Mplus* and *lcmm*. If no such constraints are required, *FLXMRglm()* can be used instead of *FLXMRglmfix()*. The final statement, *control= list(iter.max = 500, minprior = 0)* is used to set relevant model-estimation parameters. Users can and should enter values here that are most suitable for their situation. All of the user-specified settings for the LCGA (and GMM) models can of course be further adjusted to fit one's aims (see the package documentation for guidance). When the first parts of the code are run, the models are estimated and an object is generated that contains information on all estimated models (as defined by arguments *nrep* and *k*). Here, we named this object *lcgaMix*, which, when typed in the console, gives the following overview of the results:

	iter	converged	k	k0	logLik	AIC	BIC	ICL
1	2	TRUE	1	1	-1146.3234	2298.647	2311.291	2311.291
2	6	TRUE	2	2	-875.4560	1762.912	1788.200	1788.269
3	10	TRUE	3	3	-858.2982	1734.596	1772.528	1784.414

In the output we can see the fit statistics for LCGA models with 1-3 classes, based on the best log

likelihood observed across random starts. As noted above, the AIC and BIC are smallest for the 3-class model and models with more classes would have to be estimated to identify the ‘best’ LCGA model, but this would not be an optimal solution given the known data-generating model (see discussion above). For illustration, we can use the `getModel()` function to take a closer look at the 2-class model information that is contained in the `lcgaMix` object. First we get the model summary (`summary(lcga2)`), showing the individual class proportions, log likelihood value, degrees of freedom and the BIC:

```

      prior size post>0 ratio
Comp.1  0.5  250    265 0.943
Comp.2  0.5  250    255 0.980

'log Lik.' -875.456 (df=6)
AIC: 1762.912   BIC: 1788.2

```

Interestingly, the AIC values found with *flexmix* correspond to those found using other software, whereas the BIC values are different, suggesting that the latter are calculated differently in *flexmix*. Indeed, the package uses the total number of data points (subjects*repeated measures) as sample size (n) in the calculation of the BIC ($BIC = \ln(n)k - 2\ln(\text{LogLik})$), whereas the other covered packages use the number of subjects⁵. Note that in line with the BIC calculation, the presented class-sizes reflect the number of data points per class, rather than the number of subjects. Here, the number of subjects can be calculated by dividing the size-column (`lcgaMix@models$`2`@size`) by the number of repeated measurements (here 5). Note that this only works when the number of follow-up measurements does not vary across subjects. To look at the class-specific model-parameters, we can type ‘`parameters(lcga2)`’ in the console, which in this case yields the following results:

```

               Comp.1   Comp.2
coef.(Intercept) 9.9558053 10.498145
coef.time         0.2656535  1.766542
sigma             1.2149648  1.214965

```

Here, we see the class-specific intercepts and slopes (i.e. `coef.time`) and the residual term (`sigma`) that is the same across classes, as requested in the code. Unsurprisingly, the results are similar to those found using the other software packages (see **Tables 1** and **3**).

⁵ For instance, for the 2-class LCGA model, the BIC calculation used in *lcmm* is: $\ln(100)*6 - 2(-875.46)=1778.5$ and the calculation used in *flexmix* is: $\ln(500)*6 - 2(-875.46)=1788.2$.

2.2.2 GMM-1 in flexmix

To run a GMM with a random intercept using the *flexmix* package, we use the ‘`LXMRlmm()`’ function instead of the ‘`LXMRglmfix()`’ function in ‘`stepFlexmix()`’:

```
# set random number seed
set.seed(1234)

# estimate gmm models with 1-3 classes
gmmMix <- stepFlexmix(~.|ID, k = 1:3, nrep = 100, model =
  FLXMRlmm(y~ time, random = ~ 1, varFix = c(Random = FALSE,
  Residual = TRUE)), data = gmmdata, control = list(iter.max =
  1000, minprior = 0))

# inspect and compare model-fit results
gmm_Mix

# extract a specific model object from the gmmMix object.
gmm2 <- getModel(gmmMix, which = 2)

# inspect the model, for example:
summary(gmm2)
parameters(gmm2)

# calculate class sizes by dividing the provided sizes by the
# number of time points per subject, for example:
gmmMix@models$`2`@size/5
```

The general structure of the code is similar to that for the LCGA, except for the used function ‘`FLXMRlmm()`’, which enables us to use a linear mixed models function, and thus, inclusion of both fixed and random effects. The code for the GMM includes the term ‘`random=~1`’ to indicate that we want to estimate a random intercept. The element `varFix=c(Random=FALSE, Residual=TRUE)` is used to indicate if the random effect variance and residual variance should be fixed across classes or estimated for each class separately. Here, the intercept variances are estimated class-specific by stating `Random=FALSE` and the residual variance is kept constant across classes by stating `Residual=TRUE`. After running this code, the general model output, obtained by typing ‘`gmm_Mix`’ in the console, looks like this:

	iter	converged	k	k0	logLik	AIC	BIC	ICL
1	8	TRUE	1	1	-1041.9742	2091.948	2108.807	2108.807
2	24	TRUE	2	2	-855.3408	1726.682	1760.399	1760.901
3	34	TRUE	3	3	-838.7213	1701.443	1752.018	1763.516

The AIC and BIC are notably lower than those for the LCGA models, indicating that these models generally provide a better description of the data. As above, we see that the AIC and BIC do not reach a minimum for 2 classes because the model does not account for the slope variance in the model. A GMM-2 model is better suited for the data at hand. Still, if – for illustrative purposes - we select the 2-class model for further inspection by typing ‘`summary(gmm2)`’ in the console, we get the following information:

```
      prior size post>0 ratio
Comp.1 0.498  250    285 0.877
Comp.2 0.502  250    275 0.909

'log Lik.' -855.3408 (df=8)
AIC: 1726.682    BIC: 1760.399
```

We can see that the class sizes are the same as those found in the LCGA. The model parameter estimates can be inspected next by typing ‘`parameters(gmm2)`’ in the console:

```
coef.(Intercept) 10.4950739 9.9597174
coef.time        1.7696437 0.2682005
sigma2.Random    0.3883217 0.3084847
sigma2.Residual  1.1317975 1.1317975
```

The shown coefficients include the fixed intercept (`coef.(Intercept)`) and slope (`coef.time`) for each of the estimated classes. The requested class-specific intercept variances are given in the row `'sigma2.Random'`. Finally, the overall residual variance is provided (`'sigma2.Residual'`) and is fixed across classes, as requested in the code.

2.2.3 GMM-2 with flexmix

To run a GMM with random intercepts and slopes, the code needs to be slightly adjusted, by adding time to the random effects to be estimated (`random=~1+time`). All other arguments can be adjusted according to the user's needs.

```
# set random number seed
set.seed(1234)

#Run models with 1-4 classes
gmmMix2 <- stepFlexmix(~.|ID, k = 1:3, nrep = 100, model
= FLXMRlmm(y~ time, random = ~ 1 + time,
varFix=c(Random=F, Residual=T)), data = mydata,
control=list(iter.max = 1000, minprior = 0))

gmm2_2 <- getModel(gmmMix2, which = 2)
```

After running the code, we can again look at the model-fit summary by typing `'gmm_Mix2'` in the console.

```
iter converged k k0 logLik AIC BIC ICL
1 32 TRUE 1 1 -857.5481 1727.096 1752.384 1752.384
2 611 TRUE 2 2 -828.3619 1680.724 1731.299 1732.411
3 827 TRUE 3 3 -824.5532 1685.106 1760.969 1787.593
```

The models took much more iterations to converge compared to LCGA and GMM-1, but all estimations did eventually converge as shown in the second column (`'converged'`). When comparing the AIC and BIC, we can see that the lowest values are obtained for the 2-class model, indicating – not surprisingly – that this model best describes the simulated data. Further inspection of this two-class model (i.e. `'summary(gmm2_2)'`) shows us:

```
prior size post>0 ratio
Comp.1 0.501 250 330 0.758
Comp.2 0.499 250 310 0.806

'log Lik.' -828.3619 (df=12)
AIC: 1680.724 BIC: 1731.299
```

Here, we can see that both classes are of equal size, as expected and similar to the LCGA and GMM-1 results. Finally, we can take a look at the estimated model parameters (`'parameters(gmm2_2)'`), which look like this:

```
Comp.1 Comp.2
coef.(Intercept) 10.49933978 9.95371893
coef.time        1.76293406 0.26677424
sigma2.Random1   0.28890857 0.09849563
sigma2.Random2   -0.07222707 -0.04435989
sigma2.Random3   -0.07222707 -0.04435989
sigma2.Random4   0.11462830 0.10657086
sigma2.Residual  0.86561728 0.86561728
```


The intercepts (`coef.(Intercept)`) and slopes (`coef.time`) as well as the residual variance (`sigma2.residual`) are quite easily identified. The other parameters represent the estimated class-specific random-effects variance-covariance matrix, which has the following structure:

	Intercept	Slope
Intercept	<code>sigma2.Random1</code>	<code>sigma2.Random2</code>
Slope	<code>sigma2.Random3</code>	<code>sigma2.Random4</code>

The first random term (`sigma2.Random1`) represents the class-specific intercept variance. The entries (`sigma2.Random2` and `sigma2.Random3`, represent the off-diagonal elements of the matrix (i.e. the random intercept-slope covariances). These both have the same value as the matrix is symmetric. Finally, the term `sigma2.random4` represents the class-specific slope variance. The parameters are also presented in **Table 3**.

2.3 Comparison between R-packages and Mplus

2.3.1 Comparison of modeling results

As there are plenty of tutorials on the use of Mplus for LCGA and GMM (see e.g., www.statmodel.com), this section will not go into the code and output, but will merely use the results of the demonstrated LCGA, GMM-1 and GMM-2 analyses when conducted in *Mplus*. These results will be used for comparison to the R-packages. The used Mplus script is provided in **Appendix 2**.

The fit results of the analyses are shown in **Table 1** and **2**. The model fit of the LCGA and GMMs is comparable to that found with the R-packages. The number of estimated parameters per GMM model in *Mplus* is higher than in *lcmm* and the same as in *flexmix*, as the latter and *Mplus* both estimate a complete random effects covariance matrix per class. All packages showed the lowest BIC for the 2-class GMM-2 model, and thus identified the same model as the optimal choice. Some differences in the results, such as the optimal log-likelihood values for the more complex 2- and 3-class GMM-2 models could be explained by differences in the default settings, implementation of random starts and exact estimation methods across packages.

The 2-class model-parameter estimates of *Mplus* and the R-packages are shown in **Table 3**. Here, we can again see that the obtained results do not differ strongly, with some slight variations in the (non-significant) covariance estimates. Note that a stable non-zero estimate of the intercept-slope covariance was not expected, as the data were simulated to contain zero covariance between slopes and intercepts.

Table 3: Estimated coefficients for the 2-class models estimated by LCGA and GMM in simulated data

Model	Class	parameter	Software package		
			<i>lcmm</i>	<i>flexmix</i>	<i>Mplus</i>
LCGA	Class1	Intercept	10.5	10.0	10.5
		Slope	1.8	0.3	1.8
	Class2	Intercept	10.0	10.5	10.0
		Slope	0.3	1.8	0.3
GMM-1	Class1	Intercept	10.0	10.5	10.0
		Intercept variance	0.3	0.4	0.3
		Slope	0.3	0.3	0.3
	Class2	Intercept	10.5	10.0	10.5
		Intercept variance	0.4	0.3	0.4
		Slope	1.8	1.8	1.8
GMM-2	Class1	Intercept	10.5	10.5	10.0
		Intercept variance ¹	0.2	0.3	0.1
		Slope	1.8	1.8	0.3
		Slope variance ¹	0.1	0.1	0.1
		Slope-intercept covariance ¹	-0.07	-0.07	-0.05
	Class 2	Intercept	10.0	10.0	10.5
		Intercept variance	0.2	0.1	0.3
		Slope	0.3	0.3	1.8
		Slope variance	0.1	0.1	0.1
		Slope-intercept covariance	-0.06	-0.04	-0.08

¹Calculated using the squared proportionality coefficient

2.3.2 Comparison of modeling approaches

Looking at the code examples and the results obtained with the different packages based on the same data, we can see that the same general results are obtained with each of the covered packages and that the R-packages yield results that are similar to those found in *Mplus*. However, each package takes a slightly different approach to the configuration of the estimated models, the estimation, and the calculation of the indices of fit, which can lead to slightly different results. These differences may be more pronounced when working with more complex models and real data and should therefore be kept in mind. Some, but not all, differences that should be kept in mind are:

- The model residuals can be estimated in different ways when doing LCGA or GMM and this is reflected in the packages' default settings. In *lcmm*, a single residual is estimated for the complete model. The default in *flexmix* (FLXMRglmfix and FLXMRlmm) is to estimate residuals for each class. In *Mplus*, the default is to estimate a single residual for each time point but not for each class separately. In each package, the estimation of the residuals can be adjusted to the user's need.
- The estimated log likelihood values are similar across packages, but the calculation of the indices of fit differed slightly. *Lcmm* and *Mplus* compute the BIC using the total number of subjects for n . The used version of *Flexmix* (v2.3-15) uses the total number of measurements (i.e. subjects*time points) for n , leading to comparatively higher BIC values. In either case, it is easy to (re)calculate the BIC with a different n if this is required.
- The class-specific random effects co-variance matrices are estimated differently by *lcmm* compared to *flexmix* and *Mplus*. *Lcmm* estimates a single matrix for the highest numbered class and yields proportional coefficients that can be used to calculate the variance-covariance matrices for the other classes. *Flexmix* and *Mplus* estimate the full matrix for each class. This makes the latter models more complex and computation times can be longer for the latter packages than for *lcmm*. An interesting recent paper by McNeish et al. (2022) delves deeper into the distinction between the two approaches (proportional vs. unconstrained GMM).
- Although all three packages allow for the use of random starts, the way this is implemented is does not seem to be completely similar across packages. In addition, *Mplus* is the only package to display all and not a single top 'final stage log likelihood values at local maxima' by default and automatically notify the user if the best log likelihood value is replicated across random starts or not. This notification provides an additional source of information to judge whether the estimated model is on a local or global maximum. Hipp and Bauer (2006) do (tentatively) recommend checking this, but this information does not seem to be provided by default by *lcmm* nor *flexmix* (as far as can be judged from the package documentation).

3. Final remarks

The current tutorial presented relatively straightforward LCGA and GMM examples to get you started. In many cases, additional steps or model elements are required, such as the fitting of non-linear growth functions or the use of alternative distributions and/or link functions (e.g., when working with count or binary outcomes). These topics are outside the scope of the current tutorial. Some subjects of interest are briefly discussed below.

3.1 Missing data

In theory, LCGA and GMM have the advantage that they can be estimated using full information maximum likelihood (FIML). This means that, assuming that the missing data are missing at random (MAR; Collins et al., 2001), both subjects with complete and with incomplete data on the modeled dependent variable can be included in the analysis and the model will be estimated based on the available data. This approach has been fully implemented in *Mplus*. We are not entirely up to date on the exact implementation of FIML in the covered R-packages and also have not explored it in this tutorial. It seems that *lcmm* can also be run with missing values on dependent variables (e.g., see: <https://github.com/CecileProust-Lima/lcmm/issues/70>). For *flexmix*, we could not find explicit information on implementation of FIML to handle missing data on the dependent variable(s), so this is something that users may want to look into before deciding on whether to use the package.

Missing values on variables that are part of the study (e.g., covariates) but not explicitly modeled

as outcomes in an LCGA or GMM model are not always handled by FIML (e.g., in *lcmm*), but can be handled using multiple imputation (e.g., Schafer & Graham, 2002), e.g., using the package ‘mice’ (van Buuren, & Groothuis-Oudshoorn, 2011). In terms of implementation, this requires more steps than FIML. An excellent introduction to multiple imputation can be found here: <https://stefvanbuuren.name/fimd/>.

3.2 Covariates

It is possible to add covariates to LCGA and GMM models. In *lcmm*, covariates with both fixed or random effects in the growth model can be added as part of the model formula arguments in `hlme()`. In addition, the `classmb` argument can be used to add covariates of class membership. In *flexmix*, covariates can be included with fixed or random effects as well. Adding covariates of class-membership may be harder than for *lcmm*, as there seems to be no dedicated argument (like ‘classmb’ in *lcmm*) to do this.

3.3 Non-linear growth

The demonstrated examples can be extended to fit mixture models based on non-linear growth. If such growth is plausible, the growth model that underlies the LCGA or GMM can be extended with a non-linear term (e.g., quadratic or cubic), or can be changed to e.g., a logarithmic function. Addition of such terms is easily done in all covered packages, although the exact implementation may vary.

3.4 Other R packages

For the current tutorial, only two R-packages were selected. However, before selecting *lcmm* and *flexmix* more packages were reviewed. Eventually, the latter two were selected mainly because they contain functions that make it relatively straightforward to write the code for estimating latent classes based on linear mixed growth models. This makes them relatively user-friendly and accessible for applied researchers. However, other R-packages exist that can be used to conduct the same kinds of analyses (and much more). ‘*OpenMX*’ (Neale et al., 2016) is an excellent all-round SEM package that can also be used to run GMM. The package ‘*mixAK*’ (Komárek & Komárková, 2014) allows for (Bayesian) estimation of mixtures based on generalized linear mixed models and can also be used for GMM and related analyses, including those with multivariate outcome data. The package ‘*mixtools*’ (Benaglia et al., 2009) is a general package for finite mixture modeling, with the added option to estimate latent classes in non-parametric settings and based on regression models with various link functions. In a recent paper van der Nest et al. (2021) provide a comprehensive and up to date overview of available software.

References

- Benaglia T, Chauveau D, Hunter DR, Young D. (2009). *mixtools*: An R Package for Analyzing Finite Mixture Models. *Journal of Statistical Software* 32(6): 1–29.
- Collins LM, Schafer JL et al. (2001). A comparison of inclusive and restrictive strategies in modern missing data procedures. *Psychological Methods* 6(4): 330–351.
- Curran PJ, Obeidat K, Losardo D. (2010). Twelve Frequently Asked Questions About Growth Curve Modeling. *Journal of Cognitive Development* 11(2): 121–136.
- Hipp JR, Bauer DJ. (2006). Local solutions in the estimation of growth mixture models. *Psychological Methods* 11(1): 36–53.
- Jones BL, Nagin DS, Roeder K. (2001). A SAS procedure based on mixture models for estimating developmental trajectories. *Sociological Methods & Research*, 29: 374–393.
- Jung T, Wickrama, KAS. (2008). An Introduction to Latent Class Growth Analysis and Growth Mixture Modeling. *Social and Personality Psychology Compass* 2/1: 302–317.
- Komárek A, Komárková L (2014). Capabilities of R Package *mixAK* for Clustering Based on Multivariate Continuous and Discrete Longitudinal Data. *Journal of Statistical Software* 59(12): 1–38.
- Lazarsfeld PF (1950). The interpretation and computation of some latent structures. In S. A. Stouffer et al., *Measurement and prediction*. Princeton: Princeton Univ. Press, Ch. 11.
- Lazarsfeld PF, Henry P (1968). *Latent structure analysis*. New York: Houghton Mifflin.
- Leisch F (2004). *FlexMix*: A general framework for finite mixture models and latent class regression in R. *Journal of Statistical Software* 11(8): 1–18.
- McNeish D, Harring JR, Bauer DJ. (2022). Nonconvergence, covariance constraints, and class enumeration in growth mixture models. *Psychological Methods*.
<https://doi.org/10.1037/met0000456>
- Mislevy RJ, Verhelst N (1990). Modeling item responses when different subjects employ different

- solution strategies. *Psychometrika* 55: 195–215
- Muthén LK, Muthén BO (2013). *Mplus: Statistical analysis with latent variables. User's guide*. Los Angeles, CA: Muthén & Muthén
- Neale MC, Hunter MD, Pritikin JN et al. (2016). OpenMx 2.0: Extended Structural Equation and Statistical Modeling. *Psychometrika* 81(2): 535–49.
- Proust-Lima C, Philipps V, Lique B. (2017). Estimation of extended mixed models using latent classes and latent processes: the R package lcmm. *Journal of Statistical Software* 78:1–56
- Schafer JL, Graham, J. (2002) Missing Data: Our View of the State of the Art. *Psychological Methods* 7: 147–177.
- Rabe-Hesketh S, Skrondal A, Pickles A. (2004). GLLAMM manual. U.C. Berkeley Division of Biostatistics Working Paper Series. Working Paper 160.
- 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: 565–576
- van Buuren S, Groothuis-Oudshoorn K. (2011). MICE: Multivariate Imputation by Chained Equations in R. *Journal of Statistical Software*, 45(3): 1–67.
- van de Schoot R. (2015). Latent trajectory studies: the basics, how to interpret the results, and what to report. *European Journal of Psychotraumatology* 6: 27514.
- van de Schoot R, Sijbrandij M., Winter SD., Depaoli S, Vermunt JK. (2017). The GRoLTS-checklist: Guidelines for reporting on latent trajectory studies. *Structural Equation Modeling*, 24, 451–467.
- van der Nest G, Lima Passos V, Candel MJJM, et al. (2020). An overview of mixture modeling for latent evolutions in longitudinal data: Modeling approaches, fit statistics, and software. *Advances in Life Course Research* 43: 1–17.
- Vermunt JK, Magidson J. (2016). *Technical guide for latent GOLD 5.1: Basic, advanced, and syntax*. Belmont, CA: Statistical Innovations.
-

Appendix 1: Simulation code

```
#Download and open package MASS.
```

```
library(MASS)
```

```
#Simple function to simulate longitudinal data with random  
intercept(s) and slope(s).
```

```
ml_data <- function(n_pers, n_time, beta_int, beta_slo_time, beta_slo_covar,  
beta_slo_interact, mean_i, var_i, mean_s, var_s, cov_is, mean_r, var_r) {  
  
  REff = mvrnorm(n_pers, mu=c(mean_i,mean_s), Sigma=rbind(c(var_i, cov_is),  
                                                           c(cov_is, var_s) ))
```

```
  colnames(REff) = c("intercept","slope_time");
```

```
  dat= data.frame(ID = rep(1:n_pers, each=n_time),  
                  time = rep(1:n_time, times=n_pers),  
                  covar= c(rep(0,n_pers/2, each=n_time),rep(1,n_pers/2,  
                  each=n_time)),  
                  int = rep(REff[,1], each=n_time),  
                  slo = rep(REff[,2], each=n_time),  
                  slo_cov=rep(beta_slo_covar, each=n_time),  
                  slo_interact=rep(beta_slo_interact, each=n_time),  
                  y = NA)
```

```
  dat$time <- dat$time-1
```

```
  dat$interact <- dat$time*dat$covar
```

```
  y = with(dat, (beta_int + int) + (beta_slo_covar)*covar  
            + (beta_slo_time + slo)*time  
            + (beta_slo_interact)*interact  
            + rnorm(n=n_pers*n_time, mean=mean_r, sd=sqrt(var_r)))
```

```
  dat$y <-y
```

```
  return(dat) }
```

```
# simulate tutorial data
```

```
set.seed(2002)
```

```
d1 <- ml_data(  
n_pers=100,  
n_time=5,  
beta_int=0,  
beta_slo_time=0.3,  
beta_slo_covar=0.5,  
beta_slo_interact=1.5,  
mean_i=10,  
var_i=0.13,  
mean_s=0,  
var_s=0.09,  
cov_is=0,  
mean_r=0,  
var_r=1)
```

```
# trim the number of variables
```

```
mydata <- d1[,c(1,2,3,8)]
```

```
# plot the data (Figure 1)
```

```
Library(ggplot2)
```

```
mydata$covar <- as.factor(mydata$covar)
```

```
plot1 <-ggplot()+geom_line(aes(y=y, x=time, group=ID, colour=covar),  
data=mydata, show.legend = FALSE)
```

```
plot1 <- plot1 + theme(legend.position = "none") + theme_classic()
```

```
plot1
```

Appendix 2: used Mplus code

All code shown for the 2-class models. Note that each model was run with 1, 2 and 3 classes (by changing the `classes=c()` statement and by adding/removing the relevant class-specific parameters (i.e. `%c#1%`, `%c#2%` and `%c#3%`)

LCGA

```
Data: file="mydata_wide.dat";

Variable:
! provide column names and indicate which
! variables to include in analyses

names are rownr ID c y1 y2 y3 y4 y5;
usevariables are y1 y2 y3 y4 y5;

! State how many classes to estimate
classes=c(2);

! provide analysis details
Analysis:
type=mixture;
starts=100 20;
estimator=MLR;

! Define the overall growth model and
! determine which parameters are estimated !
overall and which are estimated for each ! class
specifically.
Model:
%overall%
! define growth model
i s | y1@0 y2@1 y3@2 y4@3 y5@4;
! fix intercept and slope variances to
! zero:
i@0;
s@0;
! constrain residuals to be equal across
! time points:
y1-y5 (1);

! For each of the classes request the mean (fixed)
intercept and slope to be estimated.
%c#1%
[i-s];

%c#2%
[i-s];
```

GMM-1

```
Model:
%overall%
i s | y1@0 y2@1 y3@2 y4@3 y5@4;
s@0;
y1-y5 (1);

%c#1%
[i-s];
i;

%c#2%
[i-s];
i;
```

GMM-2

```
Model:
%overall%
i s | y1@0 y2@1 y3@2 y4@3 y5@4;
y1-y5 (1);

%c#1%
[i-s];
i-s;
i with s;
```

```
%C#2%
[i-s];
i-s;
i with s;
```