

SAA for many N level multilevel models

Welcome to an SAA for fitting many model types developed for Stat-JR v1.0.5

Input questions

Firstly on this page you will need to specify the dataset required from the list of available datasets.

Which dataset do you wish to use:

Submit

Next you need to choose many options including the response, estimation method, clustering variables and predictor variables (both continuous and categorical) from the chosen dataset. After choosing these variables the SAA will run and you will see a block of text describing how many observations are to be used at the bottom of this page. The rest of the analysis will appear in pages 2-12.

What estimation method do you want to use:

IGLS

What is the response variable:

normexam

What distribution are you going to assume:

Normal

Please enter your possible (nested) classifications / levels (lowest first, not including level-1):

school

Are there any continuous predictors that need including in all models:

No

Are there any categorical predictors that need including in all models:	No
Do you want to include any continuous predictors as candidates for inclusion in the models:	Yes
Which continuous predictors do you want to consider:	standlrt,avslrt
Do you want to include any categorical predictors as candidates for inclusion in the models:	Yes
Which categorical predictors do you want to consider:	girl,schgend,vrband
What selection type do you require:	Forward pass
Do you want to test for random slopes:	No
Do you want to test for interactions:	No
How do you wish to compare models:	Likelihood Ratio

The Analysis Assistant you are currently using is designed to work on complete datasets only and so as a pre-processing step we have to remove any rows that contain missing data in columns used in the analysis that follows. For now the list of columns to be considered is: normexam, school, standlrt, avslrt, girl, schgend, vrband. There are 0 (0.0%) rows that get deleted This results in a dataset of 4059 rows.

On the next page we will look at the shape of the response and, in the case of normal responses, decide whether to log transform.


Exploring the response

We will begin our analysis of the dataset by doing some basic data exploration.

You have chosen normexam as your response variable and so a first step is to take a look at this variable and assess its suitability for a normal model. The summary statistics for the variable are in the table below:

Observations	4059
Mean	0.0
Standard Deviation	0.999
Median	0.004

We also look at a histogram of normexam to see if it is approximately normally distributed. Although in modelling the response in terms of a set of predictors it is what is unexplained (the model residuals) that need to be normally distributed, it is still useful to look at the response variable as a very skewed variable will often lead to very skewed residuals.



Here the distribution is reasonably symmetric with skewness value 0.004.

There are no obvious outliers in normexam.

**Do you want to log transform the
response variable?:**

No

Exploring the predictors individually


We can also look at each of the predictor variables in turn in isolation.

For categorical predictors we are looking at how common each category is in the dataset. In particular we are checking for rare categories which might cause difficulties in modelling and might therefore be usefully merged with other categories (though this would need to be done outside this SAA).

For predictor girl we see the following:

girl	N	Percentage
0	1623	39.985
1	2436	60.015
Total	4059	100


None of the categories of girl have fewer than 5 observations.



For predictor schgend we see the following:

schgend	N	Percentage
1	2169	53.437
2	513	12.639
3	1377	33.925
Total	4059	100


None of the categories of schgend have fewer than 5 observations.



For predictor vrband we see the following:

vrband	N	Percentage
1	1176	28.973
2	2344	57.748
3	539	13.279
Total	4059	100


None of the categories of vrband have fewer than 5 observations.



For continuous predictors we are interested in looking at summary statistics, the shape of the distribution and any unusual values. If the distribution is skewed then we might want to transform the variable before fitting it in the model although it is more important to consider transformations of the response variable and remember what is important is whether the relationship between the response and predictor is linear. If there are unusual values we will want to check that the unusual values are correct and not errors and also whether we may want to treat the variable differently. Another possibility for unusual shaped distributions is to instead categorise the variable into ranges of values.

For predictor standlrt we see the following:

Name	standlrt
Observations	4059
Mean	0.002
Standard Deviation	0.993
Median	0.04




Here the median is larger than the mean and there is significant skew to the left. The skewness value is -0.128 . Here the statistical significance may be to some degree due to the large sample size as from a practical perspective values of skew less than 2 in absolute magnitude are not considered too big a skew.

There are no obvious outliers in standlrt.

For predictor avslrt we see the following:

Name	avslrt
Observations	4059
Mean	0.002
Standard Deviation	0.315
Median	-0.02



Here the median is larger than the mean and there is significant skew to the left. The skewness value is -0.203 . Here the statistical significance may be to some degree due to the large sample size as from a practical perspective values of skew less than 2 in absolute magnitude are not considered too big a skew.

There are no obvious outliers in avslrt.

Assessing the relationship between the response and individual predictors

Once we are happy with our response variable and our set of predictors we now want to have a preliminary look at them together before progressing to the univariable modelling.

For the categorical predictors it is worth looking at the mean value of the response in each category to assess if there are differences. We can then formally test this with a t-test for binary predictors or an ANOVA for predictors with more than 2 categories.

Here is a tabulation of the response, normexam for predictor girl with category 1 having the largest mean and category 0 the smallest.

Category	N	Mean	Standard Deviation	Median
0	1623	-0.14	1.025	-0.129
1	2436	0.0933	0.97	0.0735

The formal test is as follows:

There are two groups in the data:

The first group has 1623 observations with mean -0.14 standard deviation 1.026.

The second group has 2436 observations with mean 0.093 standard deviation 0.97.

We are trying to test a hypothesis as to whether the two groups differ in their (population) means by a statistically significant amount. Statistical significance is related to how likely a result is to be a chance occurrence. Here we are trying to differentiate between a real difference (no matter how small) and a difference that may have occurred due to the samples we have chosen.

The mean difference is 0.234 with the second group having the larger sample mean.

We need to quantify if this difference is large relative to the variability in the data. To do this we calculate the standard error of the difference. This is a function of the variabilities in the samples from group A and group B combined with their sample sizes. The bigger the 2 variabilities the larger the standard error, whilst the smaller the variability the smaller the standard error.

For our data the standard error of the mean difference is 0.032 and we divide our observed difference by this standard error to give a test statistic with value 7.266.

This test statistic is then compared to a t distribution with degrees of freedom equal to the sum of the sample sizes in each group (4059) - 2. In this case a t distribution with 4057. This t table has values of 1.961 for $p=0.05$ and 2.577 for $p=0.01$.


As $7.266 > 2.577$ our p value is less than 0.01 and we have strong evidence to reject the null hypothesis (at the $p=0.01$ level).

The p-value is in fact less than 0.0001.

The t test assumes that the distribution of the response in each group follows a Normal distribution. We could check this by looking at histograms of the variable in each group. If we were concerned about the normality assumption then we could instead use a Mann Whitney (MW) test.

A Mann Whitney test works simply on the order (or ranks) of the responses across the two groups. So the response variable is firstly sorted and then each value is ranked. The ranks for each group are then summed and the value that is larger is compared with what would be expected if there was no difference between the groups.

In this case the MW U statistic is 1740511 which for samples of size 1623 and 2436 corresponds to a p value of less than 0.0001.



Here is a tabulation of the response, normexam for predictor schgend with category 3 having the largest mean and category 1 the smallest.

Category	N	Mean	Standard Deviation	Median
1	2169	-0.0984	0.984	-0.129
2	513	0.0234	1.055	0.0735
3	1377	0.146	0.982	0.134

The formal test is as follows:


	df	SS	MS	F
Between groups	2	50.66	25.33	25.69
Within groups	4056	3999.0	0.986	
Total	4058	4049.0	0.998	
Pooled within-group S.D.	0.993			
Between-group variance component	0.0206			

For the ANOVA we are testing whether there are differences in the means of the response variable between the different groups. As shown in the table above this is done by constructing an ANOVA table that compares how much of the variability in the data is within the groups compared to between the groups. This results in a test statistic that follows an F distribution with 2 and 4056 degrees of freedom. This F table has values of 3.692 for $p=0.05$ and 5.305 for $p=0.01$.

As $25.694 > 5.305$ our p value is less than 0.01 and we have strong evidence to reject the null hypothesis (at the $p=0.01$ level).

The p-value is in fact less than 0.0001.

Category	N	Mean	S.E.M.
1	2169	-0.0984	0.0213
2	513	0.0234	0.0438
3	1377	0.146	0.0268



Here is a tabulation of the response, normexam for predictor vrband with category 1 having the largest mean and category 3 the smallest.

Category	N	Mean	Standard Deviation	Median
1	1176	0.736	0.856	0.747
2	2344	-0.142	0.827	-0.129
3	539	-0.99	0.831	-1.029

The formal test is as follows:


	df	SS	MS	F
Between groups	2	1212.0	605.8	865.8
Within groups	4056	2838.0	0.7	
Total	4058	4049.0	0.998	
Pooled within-group S.D.	0.836			
Between-group variance component	0.528			

For the ANOVA we are testing whether there are differences in the means of the response variable between the different groups. As shown in the table above this is done by constructing an ANOVA table that compares how much of the variability in the data is within the groups compared to between the groups. This results in a test statistic that follows an F distribution with 2 and 4056 degrees of freedom. This F table has values of 3.692 for $p=0.05$ and 5.305 for $p=0.01$.

As $865.766 > 5.305$ our p value is less than 0.01 and we have strong evidence to reject the null hypothesis (at the $p=0.01$ level).

The p-value is in fact less than 0.0001.

Category	N	Mean	S.E.M.
1	1176	0.736	0.0244
2	2344	-0.142	0.0173
3	539	-0.99	0.036




For the continuous predictors we can look at correlations with the response and scatterplots to see if there is a linear relationship.

Predictor : standlrt

The Pearson correlation between normexam and standlrt is 0.592 (p value < 0.001).

The Spearman rank correlation between normexam and standlrt is 0.58 (p value < 0.001).




The graph includes best fitting curves for a constant, linear, quadratic and cubic relationship between normexam and standlrt. In this case a quadratic relationship is most appropriate and you might consider including a squared term in the predictor list.

Predictor : avslrt

The Pearson correlation between normexam and avslrt is 0.288 (p value < 0.001).

The Spearman rank correlation between normexam and avslrt is 0.273 (p value < 0.001).



The graph includes best fitting curves for a constant, linear, quadratic and cubic relationship between normexam and avslrt. In this case a quadratic relationship is most appropriate and you might consider including a squared term in the predictor list.

Choosing appropriate random classifications

We begin this section by deciding which of the possible random classifications to include in the modelling.

This is done by fitting combinations in turn and picking more complicated models if they make a significant improvement via a LR test. All models are displayed along with their likelihood in the table below:

Higher-level classifications	Deviance	Likelihood Ratio	p value
None	11509.36	-	-
school	11010.65	498.72	< 0.001

The best model based on the Likelihood has levels: school

As this is a multilevel modelling SAA we will also want to look at how the response is distributed across the levels of the model.

For this we will use the best model chosen above and look at how the variance is distributed across levels.

Variable	Coefficient	SE
Intercept	-0.0132	0.0536
school Variance	0.169	0.0324
Level 1 Variance	0.848	0.019

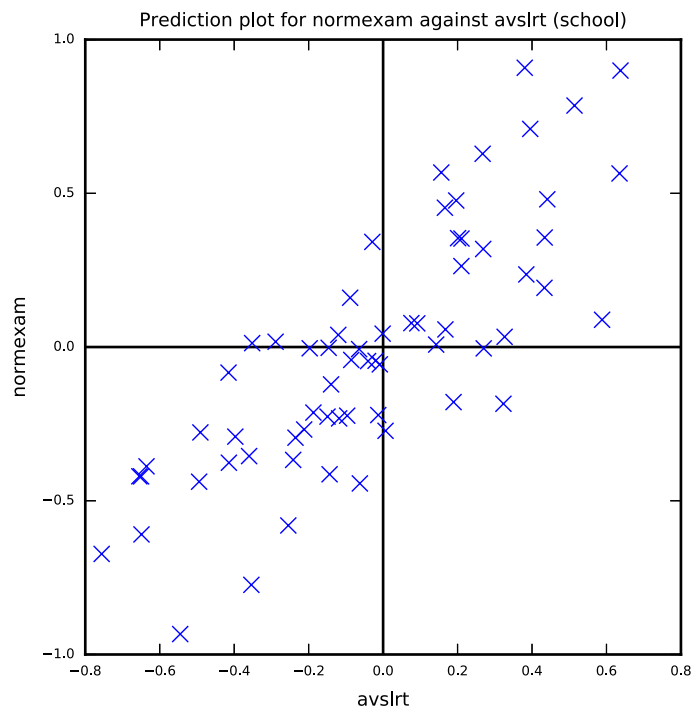

Here we see that the VPC for school = $0.169/1.016 = 0.166$, so we see that school effects explain 16.59% of the variability in normexam.



Performing univariable modelling


Our next step in modelling now that we have a set of potential predictors is to consider models for each predictor in turn along with a random intercept at each chosen classification from the best model in the last section. In the fixed part these models simply contain an intercept and the particular predictor and so for continuous predictors will be multilevel linear regressions and for categorical predictors will be multilevel generalisations of ANOVAs. In the table below we summarise the modelling by showing the coefficients for each predictor along with the p value comparing the model with that predictor with a Null model. This Univariable modelling step will identify a set of candidate predictors to be taken forward into the next stage of modelling.

Variable	Coefficient	SE	p value	Significance
standlrt	0.563	0.0125	< 0.001	***
avslrt	0.926	0.109	< 0.001	***
girl_1	0.262	0.0403	< 0.001	***
schgend_2	0.0644	0.149	0.095	N/S
schgend_3	0.258	0.117		
vrband_2	-0.82	0.0285	< 0.001	***
vrband_3	-1.614	0.042		

Which predictors we consider for the next stage of analysis will depend on their significance in the above table (but may in practice also depend on the size the effect and substantive interest of the variable though this is hard to automate). so the predictors to carry forward are: girl, avslrt, standlrt, and vrband.







Looking at correlations between predictors

Our next step is to check that none of the correlations between the predictor variables are too great as this could cause estimation problems when we add the predictors to the model together. To do this we look at all correlations between the predictor variables that have been identified as significant univariably and are thus candidates to be added to the model.

The correlations are as follows:

Variables	Correlation
(avslrt, standlrt)	0.317
(girl_1, standlrt)	0.053
(girl_1, avslrt)	0.041
(schgend_2, standlrt)	0.0
(schgend_2, avslrt)	0.001
(schgend_2, girl_1)	-0.466
(schgend_3, standlrt)	0.006
(schgend_3, avslrt)	0.02
(schgend_3, girl_1)	0.585
(schgend_3, schgend_2)	-0.273
(vrband_2, standlrt)	-0.18
(vrband_2, avslrt)	-0.06
(vrband_2, girl_1)	0.015
(vrband_2, schgend_2)	-0.039
(vrband_2, schgend_3)	0.03
(vrband_3, standlrt)	-0.525
(vrband_3, avslrt)	-0.187
(vrband_3, girl_1)	-0.053
(vrband_3, schgend_2)	0.028
(vrband_3, schgend_3)	-0.032
(vrband_3, vrband_2)	-0.457

Correlations greater than 0.8 (in magnitude) are worth looking at as they may result in model fitting problems when both predictors are included.

Performing multivariable model selection - random intercept models

In this next stage we will look at the best random intercepts model using only main effects for the variables to be considered. You have chosen to perform forward pass which is a quicker method than full forward selection. It may therefore not explore as many possible models. The predictor variables are considered in turn based on their significance in the univariable analysis and each is added to the current model. If the resulting model is a significant improvement then the predictor is kept in the model otherwise it is removed. Attention then moves on to the next predictor until all predictors are considered.

You have chosen to use Likelihood ratio tests to compare models and here change in deviance will be used to indicate whether a model is better or not by comparison with an appropriate chi-squared distribution. This method is slightly slower than the alternative Wald test which we offer as for each stage a model is compared with all its submodels (with 1 predictor removed) to work out p values for each predictor.

The most significant predictor in the univariable analysis was vrband so our starting point in multivariable modelling is the model:

$$\text{normexam}_i = \beta_0 \text{vrband_2}_i + \beta_1 \text{vrband_3}_i + \beta_2 \text{intercept}_i + u_{0, \text{school}_i}^{(2)} + e_i$$

Variable	Coefficient	SE	p value	Significance
vrband_2	-0.82	0.0285	< 0.001	*** (df=2)
vrband_3	-1.614	0.042		
Intercept	0.692	0.0462		
Between school Variance	0.0999	0.0196		
Level 1 Variance	0.607	0.0136		

Adding variable vrband is significant and so is retained in the model.

Our next step is to consider adding variable standlrt to the current model.

$$\text{normexam}_i = \beta_0 \text{vrband_2}_i + \beta_1 \text{vrband_3}_i + \beta_2 \text{standlrt}_i + \beta_3 \text{intercept}_i + u_{0, \text{school}_i}^{(2)} + e_i$$

Variable	Coefficient	SE	p value	Significance
vrband_2	-0.417	0.0319	< 0.001	*** (df=2)
vrband_3	-0.765	0.0537		
standlrt	0.391	0.0168	< 0.001	*** (df=1)
Intercept	0.348	0.0456		
Between school Variance	0.0865	0.0169		
Level 1 Variance	0.536	0.012		

Adding variable standlrt is significant and so is retained in the model.

Our next step is to consider adding variable avslrt to the current model.

$$\text{normexam}_i = \beta_0 \text{vrband_2}_i + \beta_1 \text{vrband_3}_i + \beta_2 \text{standlrt}_i + \beta_3 \text{avslrt}_i + \beta_4 \text{intercept}_i + u_{0, \text{school}_i}^{(2)} + e_i$$

Variable	Coefficient	SE	p value	Significance
vrband_2	-0.416	0.0319	< 0.001	*** (df=2)
vrband_3	-0.763	0.0537		
standlrt	0.387	0.0168	< 0.001	*** (df=1)
avslrt	0.33	0.108	0.003	** (df=1)
Intercept	0.356	0.0432		
Between school Variance	0.0728	0.0146		
Level 1 Variance	0.536	0.012		

Adding variable avslrt is significant and so is retained in the model.

Our next step is to consider adding variable girl to the current model.

$$\text{normexam}_i = \beta_0 \text{vrband_2}_i + \beta_1 \text{vrband_3}_i + \beta_2 \text{standlrt}_i + \beta_3 \text{avslrt}_i + \beta_4 \text{girl_1}_i + \beta_5 \text{intercept}_i + u_{0, \text{school}_i}^{(2)} + e_i$$

Variable	Coefficient	SE	p value	Significance
vrband_2	-0.415	0.0318	< 0.001	*** (df=2)
vrband_3	-0.758	0.0536		
standlrt	0.385	0.0168	< 0.001	*** (df=1)
avslrt	0.316	0.106	0.004	** (df=1)
girl_1	0.161	0.0317	< 0.001	*** (df=1)
Intercept	0.263	0.0465		
Between school Variance	0.0699	0.014		
Level 1 Variance	0.533	0.0119		

Adding variable girl is significant and so is retained in the model.

Our next step is to consider adding variable schgend to the current model.

$$\text{normexam}_i = \beta_0 \text{vrband}_2_i + \beta_1 \text{vrband}_3_i + \beta_2 \text{standlrt}_i + \beta_3 \text{avslrt}_i + \beta_4 \text{girl}_1_i + \beta_5 \text{schgend}_2_i + \beta_6 \text{schgend}_3_i + \beta_7 \text{intercept}_i + u_{0, \text{school}_i}^{(2)} + e_i$$

Variable	Coefficient	SE	p value	Significance
vrband_2	-0.415	0.0318	< 0.001	*** (df=2)
vrband_3	-0.759	0.0535		
standlrt	0.385	0.0168	< 0.001	*** (df=1)
avslrt	0.321	0.102	0.002	** (df=1)
girl_1	0.157	0.0331	< 0.001	*** (df=1)
schgend_2	0.191	0.0993	0.054	N/S (df=2)
schgend_3	0.152	0.0782		
Intercept	0.189	0.0543		
Between school Variance	0.063	0.0128		
Level 1 Variance	0.533	0.0119		

Adding variable schgend did not significantly improve the model, so we remove it from the model.

We have considered all variables so now run our final random intercepts model.

$$\text{normexam}_i = \beta_0 \text{vrband_2}_i + \beta_1 \text{vrband_3}_i + \beta_2 \text{standlrt}_i + \beta_3 \text{avslrt}_i + \beta_4 \text{girl_1}_i + \beta_5 \text{intercept}_i + u_{0, \text{school}_i}^{(2)} + e_i$$

Variable	Coefficient	SE	p value	Significance
vrband_2	-0.415	0.0318	< 0.001	*** (df=2)
vrband_3	-0.758	0.0536		
standlrt	0.385	0.0168	< 0.001	*** (df=1)
avslrt	0.316	0.106	0.004	** (df=1)
girl_1	0.161	0.0317	< 0.001	*** (df=1)
Intercept	0.263	0.0465		
Between school Variance	0.0699	0.014		
Level 1 Variance	0.533	0.0119		

This is our final model.

Choosing interactions

You have chosen not to investigate interactions and so this page is empty.


Adding random slopes

You have chosen not to look at random slopes and so this page is blank.

Analysing the residuals


Here we look at the residuals from the model and plot them in various ways.

We start with level 1 residuals:




Here the median is larger than the mean and there is significant skew to the left. The skewness value is -0.138 . Here the statistical significance may be to some degree due to the large sample size as from a practical perspective values of skew less than 2 in magnitude are not considered too big a skew.

There are no obvious outliers in the residuals.




If the residuals are fairly normally distributed then the points in this graph should be close to the red line.




Here you should consider whether there are any patterns in this plot. Ideally we would like to see similar variability of the residuals across the range of fitted values.

Next the level 2 residuals for intercept:



Here the distribution is reasonably symmetric with skewness value -0.17 .

There are no obvious outliers in the residuals.



If the residuals are fairly normally distributed then the points in this graph should be close to the red line.

Looking at predictions

Having fitted a model with several predictors we might like to represent this model graphically. This is more difficult than when we have only one predictor and so for now we consider each predictor in turn and set all other predictors to their mean values.

