Examining individual learning patterns using generalized linear mixed models

Sean Commins¹, Antoine Coutrot², Michael Hornberger³, Hugo J Spiers⁴, Rafael De Andrade Moral⁵

¹Department of Psychology, Maynooth University, Maynooth, Co. Kildare, Ireland.

² Laboratoire d'Informatique en Image et Systèmes d'information, CNRS, Université Claude Bernard Lyon, France.

³Norwich Medical School, University of East Anglia, Norwich NR4 7TJ, United Kingdom

⁴Institute of Behavioural Neuroscience, Division of Psychology and Language Sciences, Department of Experimental Psychology, University College London, London WC1H 0AP, United Kingdom.

⁵Department of Mathematics and Statistics, Maynooth University, Maynooth, Co. Kildare, Ireland.

Corresponding Author: Sean.Commins@mu.ie

Abstract

Everyone learns differently, but individual performances are often ignored in favour of a group level analysis. Using data from four different experiments we show that generalized linear mixed models (GLMMs) and extensions can be used to examine individual learning patterns in more depth. Producing ellipsoids and cluster analyses based on predicted random effects, individual learning patterns can be identified which, in turn, can be clustered and used for comparisons across various experimental conditions or groups. We show that this analysis can handle a range of data sets including discrete, continuous, censored and non-censored, as well as different experimental conditions, different sample sizes and trial numbers. Using this approach we show that learning a face-named paired associative task produced individuals that can learn quickly with the performance of some remaining high, but with a drop off in others. Whereas other individuals show poor performance throughout the learning period. We see this more clearly when we examined learning in a virtual navigation spatial learning task (NavWell). Two prominent clusters of learning emerged, one showing individuals that produced a rapid learning performance and a second group that learned the task slowly and gradually. We can also identify how individual learning patterns change across trials. Using data from another spatial learning task (SeaHero Quest), we show that performance of individuals mirror their age category. However, not all individuals perform according to their age, some show a pattern that more readily fits a younger category, while others show patterns that fit better with an older age cohort. Overall, using this analytical approach may help practitioners in education and medicine to identify those that might need extra help and attention. In addition, identifying learning patterns may allow further investigation of the underlying neural, biological, environmental and other factors of these individuals.

Key Words: Learning, GLMMs, Spatial, Individual.

Introduction

Learning is a dynamic process and fluctuates across time. Repeated examination of a task generally leads to an increased performance; however, learning rates are individualistic with some learning a task quicker than others, while others may never learn. Such variation across time and individuals is often well captured by variation around the mean, relying on repeated measures analysis of variances (ANOVAs) or mixed factorial ANOVAs when comparing group performance across time (see Barnhart et al. 2015; Bootsma et al. 2018 and Raboyeau et al. 2010 across spatial, visuo-motor, and lexical learning domains, and our own Farina et al., 2015). However, examining variation and how individuals perform a task across time is important. Given the replication crisis across many fields including psychology and the neurosciences, there is an increased emphasis on displaying all data points in a clear explicit manner (see Allen et al., 2019 for data visualisation using raincloud plots). As well as being more transparent, examination of individual performances and data may lead to a better evaluation of outliers, the comparison of performance in various settings (e.g in education, see Braithwaite et al., 2019), the examination of individual diagnosis or treatment plans (Chiang et al., 2020; Simon, 2001) and the ability to account for individual behavioural patterns (Seidler et al. 2015).

Generalized linear mixed models (GLMMs) are widely used as a method across many fields (Bolker et al., 2008; Demétrio et al., 2014). The recognition that data may not be normally distributed as well as the addition of random effects to the linear predictor are important features of GLMMs. In psychology and cognitive sciences GLMMs have also been increasingly used (Baayen et al., 2008) and the method has even been suggested to become the main analytical tool for quantitative data (Meteyard & Davies, 2020). GLMMs are very versatile and have been used to examine many psychological and cognitive constructs, such as associating grip strength with cognitive decline in older adults (Quesque et al., 2020; Chou et al., 2019), as well as in patients with varying forms of depression (Frith et al., 2018). In addition, they have been used to look at the relationship between cognitive processes and neural biomarkers, again in a variety of cohorts from young adolescents (Paulus et al., 2019) to patients suffering from dementia (McDade et al., 2018). One of the great strengths of GLMMs is their use in the modelling and the predicting of outcomes using longitudinal or repeated data. For example, Song et al. (2020) use baseline cardiovascular scores to predict cognitive decline and neural changes in the subsequent 21 years. Similarly, Mattei et al. (2019) showed that adults with type-2 diabetes using a Mediterranean diet were associated with higher cognitive scores compared to those that didn't in a 2-year follow-up examination.

GLMMs account for main and interaction differences (for fixed experimental effects, e.g. we may use mixed model ANOVAs for Gaussian GLMMs). Importantly, GLMMs are used to estimate variance components associated to random effects. Such random effects may arise from individual differences, with participants deviating from the grand mean with respect to time, location or other unknown factors. Although the inclusion of such random effects may be typically used as a control feature, it is also an ideal way to examine individual variation in performance with respect to task items or across time, such as learning. For example,

examining the effects of sleep loss on attention Cochrane et al. (2021) use a GLMM (and comparing to their own model) to specifically examine and predict individual vigilance scores from a number of sleep-related measures. Similarly, Kliegl, Wei, Dambacher, Yan, & Zhou (2011) provide a nice illustration of how individual performances can be examined using GLMMs. In this study, participants had to complete visual attention task and respond as quickly as possible to stimuli presented on a screen (there were 3 experimental conditions - response to changes in object, spatial and central fixation conditions). The authors assess the significance of main (fixed) effects for all three conditions. Importantly they also show how individual responses vary within a condition and that individual responses may correlate across some conditions and not others (random effects). For example, the authors showed that responses of individuals in the spatial condition were very variable (with some performing very well compared to others), whereas in the object condition there was limited variability across individuals. Further, participants' responses in the spatial condition strongly correlated (negatively) with responses for the central fixation condition, but did not correlate at all with the object condition. In addition, the authors show how individual responses differ and how correlations might change, depending on the model used.

Although used in the above examples, this type of extra individual-level analysis is seldom done, with many studies only including random effects in their models to control for dependence between experimental or observational units, i.e. to simply reflect design. As such, interesting insights into individual learning (or indeed patterns of learning across conditions) may be lost. Further to this and mirroring the criticisms of Meteyard & Davies (2020) to the multiplicity of approaches of GLMMs reporting in general, there is little guidance as to how such individual-level analysis should be examined. Here we offer a set of approaches and show how GLMMs can be used to examine individual learning patterns by exploring individual-level random effects combined with outlier detection and clustering methodologies. We illustrate how a general framework may be applied using different types of datasets, including strictly positive continuous responses, right-censored data, and discrete proportions.

Methods

A number of datasets and procedures were used to try and examine how the GLMMs and analysis would deal with different types of data (continuous, non-continuous, censored, non-censored), different numbers of participants and trials, as well as different learning tasks.

Behavioural Procedures

The Face-Name Pairs task was used as an example of discrete non-continuous data. The task is used to assess associative learning and memory and consists of 8 face-name paired stimuli, presented twice in a block; there were four blocks in total (see Zeineh et al. (2003)). Each face-name pair was presented in random order on screen for five seconds. The four blocks were either presented sequentially on the same day (massed condition) or one block of face-name pairs was presented each day for four days (spaced condition). After each of the four blocks,

a test trial was given. Each trial consisted of the eight faces presented once in random order, without their corresponding name. The number of correctly recalled names associated with each of the eight faces (out of eight) was used to measure learning performance across the four blocks.

The first data set was conducted in a controlled laboratory setting with 57 participants in the massed condition and 61 in the spaced condition. Both groups were well matched in terms of gender (massed: M/F = 31/30, spaced: M/F = 29/28), age (massed = 22.9 (SD = 1), spaced = 23.4 (SD = 1.2) and general IQ as measured by the National Adult Reading Test (NART; massed = 23.4 (SD = 1.6), spaced = 24.3 (SD = 1.7)).

The second experiment was conducted online through the online survey platform Qualtrics and had 358 participants in total (179 per condition, massed and spaced). Due to the uncontrolled online conditions the conditions were not well matched for gender (85 male/94 female for the massed condition and 70 male/109 female for the spaced condition) or age (with the massed group being older (mean = 29.4 (SD=13.1)) than the spaced group (mean = 24.1 (SD=8.8)).

We used both these data sets to explore how our analysis would deal with non-continuous data and differing numbers of participants.

The virtual navigation task (NavWell, Commins et al., 2020) was used as an example of continuous data that was right-censored (i.e. with a maximum limit). This task is the human equivalent of the Morris water maze task that is generally used to assess spatial learning and memory (Morris, 1981). The task consists of participants virtually navigating around an enclosed circular arena (22 virtual metres or 15.75 s to traverse its diameter) in an attempt to locate an invisible target, located somewhere on the ground. The hidden target only becomes visible once the participant traverses it. On subsequent trials, participants must try and recall this specific location and make their way to the hidden target as quickly as possible. To aid their recall, two cues (large shapes) are located on the wall of the circular arena. Each participant is given twelve trials; participants must try and find the target within 60 seconds. If the participant cannot locate the target within this time, he/she is transported to the hidden target and instructed to look around the arena and try and recall the location on the subsequent trials. For those that successfully locate the target, they are also told to look around the arena and try to recall the specific location. Time taken to reach the target for each of the 12 trials is used as the dependent measure. The lower the score reflects better spatial learning and memory. For this task 42 participants were used (M/F = 19/23) and mean age of 28.3 (SD = 14, range 19-62).

The Sea Hero Quest mobile video game (Coutrot et al., 2018; Spiers, Coutrot & Hornberger, 2021) was used to examine continuous, non-censored data that required semi-independent learning (the task contains increasing levels of difficulty). The Sea Hero Quest task was designed to measure human spatial navigation ability through gameplay. Currently, over 4 million people across 195 countries have downloaded the game and have played it. The game has a number of different features but primarily it involves participants virtually navigating a boat through a series of waterways and rivers to find a target – the goal is to find a sea

creature in a particular location and photograph it. Before setting off, participants are provided with a map that shows their current location and to where they need to navigate (i.e. the target). The task has a number of different levels; each level gets increasingly more difficult and contains more twists and turns (i.e. takes longer to complete), as well as having rivers that do not lead anywhere. There is no time limit to the task but unlocking a more difficult level is dependent on completing the previous one. The length (in virtual meters) to reach the target is used to measure spatial learning. Gaming ability and difference in touchscreen handling/proficiency of participants has been taken into account and controlled by normalising SHQ performance for the first two levels (see Coutrot et al., 2018 for details). Although the task primarily relies on spatial learning and memory, it also depends on other cognitive processes including the translation of a 2D map into a 3D game, planning of routes and the continuing monitoring of progress during the game.

The overall data set contains 3,317 participants, however we have just looked at a random subset of this ($n = 240$) with 4 age groups (18-20, 21-40, 41-60 and 61-80, $n = 60$ per cohort). Within each cohort we tried to match for gender and to have an even spread of age. The resulting data gives us M/F: 30/30, mean age = 19.47, SD = 0.5; M/F: 31/29; mean age = 30.5, SD = 5.7; M/F: 31/29, mean age = 50.48, SD = 5.4 and M/F: 33/27 mean age = 70.9, SD = 5.7 for the 4 groups, respectively.

Statistical procedures

Here we present the statistical methods used for each sample dataset. For a summary, see Table 1 that also includes the syntax used to fit the models in R (R Core Team, 2021).

Face-Name Pairs task

For this task the response variable is the number of times the individual made a correct association between a face and a name. This response is discrete and bounded between 0 (all incorrect) and 8 (all correct). Typically binomial models are used to analyse discrete proportion responses, also known as logistic regression (when the logit link is used). A normal distribution would be an inadequate assumption, since normal models assume unbounded and continuous responses.

Here we fitted a binomial GLMM. We included a different linear effect of trial (as a continuous predictor, ranging from 1 to 4) per learning condition (massed vs. spaced) in the linear predictor, which yielded one intercept and one slope per learning condition (i.e. four fixed effects). We also included individual-level random intercepts and slopes, assumed to be independent and to follow a normal distribution with mean zero and a variance to be estimated by the model (i.e. two variance components). We used a logit link function.

For the binomial distribution, the dispersion parameter is known (or fixed) and equal to 1. However, in case the variability of the data is larger than expected by the binomial model, it is possible to estimate this dispersion parameter in a quasi-likelihood approach to accommodate the extra-variability. Other approaches include the use of mixtures of

distributions, such as the beta-binomial, or the inclusion of an observational-level random effect within a GLMM framework.

Virtual navigation task

The response here is strictly positive (since it is the time taken until the individual reached the target), and an important feature is that the response is right-censored. This is because if the individuals had not found the target within 60 seconds, the task would be interrupted. Therefore, for these cases, we have the information that it would have taken more than 60 seconds for that individual to find the target, but we do not know exactly how long. This information can be incorporated in the modelling framework by using a different term in the likelihood for censored observations, based on the survival function (defined as the probability of the observation being more than 60 seconds, for this case). Again, the normal distribution would not be suitable, since although the data is continuous, it is strictly positive (i.e. time cannot be negative).

We fitted a right-censored gamma GLMM. The gamma distribution is very flexible, and accommodates different shapes of continuous, positive and right-skewed data. The gamma distribution can be parameterised to have a mean and a dispersion parameter proportional to the variance. For the mean parameter, we included a linear effect of trial (as a continuous predictor, ranging from 1 to 12), that is, an intercept and a slope, in the linear predictor as fixed effects, and random and independent intercepts and slopes per individual. Additionally, we included the linear effect of trial in the linear predictor for the dispersion parameter, which allowed us to model the changes in variability as the trials progressed. Both mean and dispersion were modelled with a log link.

Other potential approaches would include the use of different link functions for the mean parameter, such as the inverse link (which is called the canonical link function for the gamma model), the inclusion of semi-parametric terms in the linear predictor such as splines (to capture nonlinear behaviour), or the use of alternative distributions to the gamma, such as the inverse Gaussian, Weibull, log-normal, among others. There is a plethora of probability distributions that accommodate positive continuous data, with a continuously growing literature on the development of new models aimed at time-until-event and censored data.

Sea Hero Quest

For this study, the response is also strictly positive and continuous, as it is the length used to reach a target. In contrast to the virtual navigation task, the data is not censored, which means that the trials were not limited in terms of a maximum length. The behaviour of the response is non-linear, and therefore a simple linear model with an intercept and slope would not be sufficient.

Here, we fitted a gamma GLMM, but instead of a linear effect of trial, we included semi-parametric smooth functions based on b-splines with three knots across trials (again, taken as a continuous predictor, ranging from 1 to 5). This totals four estimated parameters to represent a curve. We estimated a different curve per age group, totalling 16 parameters, or fixed effects, in the linear predictor for the mean of the distribution. We also included

individual-level random effects, representing a different curve per individual. Therefore, each individual is associated to four random effects, instead of two as described in the previous examples. These effects do not have a clear meaning such as an intercept and slope, but still hold the information used to generate the non-linear smooth curves that describe each individual's learning behaviour. We also included the same fixed effects used for the mean parameter in the linear predictor for the dispersion, i.e. different smooth functions across trials per age group (also totalling 16 parameters). Both mean and dispersion were modelled with a log link.

Producing ellipsoids based on the bivariate normal distribution

For the models that estimate individual-level curves based on random intercepts and slopes, it is possible to plot the random intercepts versus the random slopes, and produce ellipsoids based on the bivariate normal distribution. By looking at, e.g. a 95% ellipsoid, one might observe individuals who fall beyond the area delimited by the ellipsoid and look further into their learning behaviour, since it may be considered "extreme", or "outlying", when compared to others. Moreover, when comparing e.g. treatment levels, different ellipsoids may be produced for different groups, and simple statistics may be computed from these ellipsoids, such as area and eccentricity indices, which may aid comparison of treatment levels.

The power of this type of analysis lies on how interpretable the random effects are. For the intercepts versus slopes case, it is a great tool to discriminate individuals according to their learning behaviour. The plot can be divided into four quadrants. Points close to the origin represent individuals that have intercepts and slopes very similar to the overall mean. If the response variable is directly proportional to learning, in a plot with random intercepts as the x-axis and random slopes as the y-axis, points further in the first and fourth quadrants represent the faster learners in the pool (larger slopes), while points in the second and third quadrants represent the slower learners (smaller slopes). Points further in the third and fourth quadrants represent learners that started at a lower level (smaller intercepts), while points further in the first and second quadrants represent learners that started at a higher level (larger intercepts).

For higher dimensions (e.g. four random effects per individual) it becomes harder to visualise, but multi-dimensional ellipsoids based on the multivariate normal distribution can still be computed. However in the example shown here (Sea Hero Quest) the coefficients do not have easily interpretable meaning, and therefore ellipsoids are not explored.

Clustering individuals based on their random effects

The individual-level random effects may then be used to carry out clustering. This may help to uncover groupings in the data based on a summarised individual-level profile, if this profile is well represented by the random effects. In this paper we use hierarchical clustering based on Euclidean distances and Ward's method, but any other clustering method may be used.

Assessing model goodness-of-fit

It is important to carry out analyses of the model residuals to check goodness-of-fit. If the model does not fit the data well, then inferences made based on the estimated parameters may be misleading. For the Face-Name Pairs task and Sea Hero Quest study, to assess whether the model fitted the data well we used half-normal plots with a simulated envelope for the Pearson residuals. These plots are such that if the data is a plausible realisation of the fitted model, most residuals will lie within the simulated envelope. For the virtual navigation task, since the model was fitted using the `gamlss` package (Rigby and Stasinopoulos, 2005) in R (R Core Team, 2021), we used worm-plots for the randomised quantile residuals. These are similar to the half-normal plots described above, however they do not require the construction of a simulated envelope, since the randomised quantile residuals should follow a $\text{uniform}(0,1)$ distribution if the model fits the data well.

Software

All analyses were carried out in R (R Core Team, 2021). Binomial GLMMs were fitted using package `lme4` (Bates et al., 2015), gamma GLMMs were fitted using package `glmmTMB` (Brooks et al., 2017) and the gamma GLMMs for censored data were fitted using packages `gamlss` and `gamlss.cens` (Rigby et al., 2018). Model goodness-of-fit was assessed using package `hnp` (Moral et al., 2017). All data and code are made available at https://github.com/rafamoral/individual_learning_GLMM.

Results

Face-Name pairs task

For the **online** Face-Name experiment, we began by assessing the significance of the variance component associated with the slopes. There was evidence that its inclusion significantly improved our model goodness-of-fit (LR = 154.67, d.f. = 1, $p < 0.0001$), which means that different individuals have significantly different learning speeds (represented in our model by the slopes of the linear predictor, or the linear effect of trial). Looking at the fixed effects, there was a significant interaction between trial and group condition (LR = 12.12, d.f. = 1, $p = 0.0005$) (Figure 1a). The slope for the spaced group (0.15) is smaller when compared to the massed group (0.40), i.e. on average the spaced condition is associated to slower learning. These slopes are interpreted in the log-odds scale, which means that in the massed group, adding a new trial increases the log-odds of correctly matching a face and a name by 0.40, which translates to being $e^{0.40} = 1.49$ times more likely to correctly matching when compared to the previous trial. For the spaced group, individuals are, on average, $e^{0.15} = 1.16$ times more likely to obtain correct matchings from one trial to the next one.

One of the perks of a mixed model is that we are able to obtain fitted curves at an individual-level without having to estimate a large number of parameters. In this particular case, estimating a linear predictor for all 358 individuals would require $358 \times 2 = 716$ parameters. Instead, we estimate an intercept and slope for each condition plus two variance components, related to random intercepts and slopes per individual, yielding four estimated parameters instead of 716, while maintaining the predictive power. This is because individual random effects are obtained through their posterior distribution given the data, instead of being individually maximised in the model likelihood. In fact, they are not even present in the likelihood function, since they are integrated out as latent, unobserved variables.

Since we now have individual-level intercepts and slopes, which are assumed to arise from normal distributions with mean zero and variances estimated by the mixed model, it becomes straightforward to rank individuals based on their random effects. This may be done looking only at the univariate distribution of random intercepts or slopes, or at their bivariate distribution. Taking the random effects as a bivariate sample from a bivariate normal distribution with a zero vector of means and a variance covariance matrix estimated by the mixed model, we produce ellipsoids at different probability levels (see Figure 1b). Although the overall eccentricity and area values for the two conditions are comparable, they do provide a good indication of spread. In this case, the massed condition has slightly more spread (area = 2.59) compared to the spaced (area = 2.44). Importantly, using this technique we can also examine individual learning patterns. We observe, for example, that individuals 54, 98 and 141 are some of the fastest learners in the massed learning condition, while individuals 116 and 174 are good learners in the spaced condition. Individuals 165, 170 (massed) and 129, 83 (spaced) are relatively weak at learning this task. Those individuals in the fourth quadrant (23, 111 (massed) and 157, 6 (spaced)) rapidly learn the task but plateau quickly. These individuals have high slope but low intercept values. In contrast, those in the second quadrant have high intercept but low slope values, with some of these individuals

starting well but getting less face-name pairs correct with additional trials (e.g. individual 152 in spaced condition).

For the **in-person** Face Name experiment we also assessed the significance of the variance component associated with the slopes. There was evidence that its inclusion significantly improved our model goodness-of-fit (LR = 7.18, d.f. = 1, $p = 0.0074$), which means that different individuals have significantly different learning speeds. Looking at the fixed effects, the interaction between trial and group condition was not significant (LR = 0.06, d.f. = 1, $p = 0.8089$). The main effect of group condition was also not significant (LR = 2.14, d.f. = 1, $p = 0.1438$), but the main effect of trial was (LR = 164.13, d.f. = 1, $p < 0.0001$) (Figure 1c). The slope for both groups was estimated to be approximately 0.70, and is also interpreted in the log-odds scale. This means that adding a new trial increases the log-odds of correctly matching a face and a name by 0.70, which translates to being $e^{0.70} = 2.01$ times more likely to correctly matching when compared to the previous trial. Although examination of the in-person experiment showed no significant difference between the conditions, we see more spread in the massed (area = 1.137) compared to the spaced condition (area = 0.998, Figure 1d). Similar exploration of individual learning patterns show that there are more individuals in the >85% area of the top right quadrant for the spaced compared to the massed condition (e.g. 26, 29, 33).

Virtual navigation task

Examination of the virtual navigation task (NavWell) demonstrated that the estimate of the variance component associated with the slope was very small (0.000006), and the likelihood-ratio test statistic was very close to zero (LR < 0.0001, d.f. = 1, $p = 1$). Therefore, the slopes seem to be very similar across all individuals. However, the effect of trial was highly significant (Figure 2a), both in the linear predictor for the mean (LR = 142.43, d.f. = 1, $p < 0.0001$) and for the variance (LR = 43.83, d.f. = 1, $p < 0.0001$). This means that the learning increases from one trial to the other, with the time taken until reaching the target decreasing by $1 - e^{-0.084} = 8\%$, on average, for the subsequent trial. Also, the variance of the data decreases from trial to trial, on average, by $1 - e^{-0.077} = 7.4\%$, until almost all participants have learned how to reach the target. Since the random slopes are based on a normal distribution with very low variance, they are all very close to zero, i.e. there is little to no individual deviation from the overall mean slope. Nevertheless, we can still rank individuals based on their predicted random intercepts. In this experiment time decreases with learning, as such individuals in the third quadrant (e.g. 29 & 17) show good learning across the trials, whereas individual 14 (first quadrant) is weaker (Figure 2b).

Even though the random slopes were very close to zero, we again carried out a hierarchical clustering analysis based on both the individual-level intercepts and slopes. Over the full 12 trials, two clusters emerged - those individuals that could be classified as fast learners and those that showed a slower learning pattern (Figures 2c and 2f). The majority of individuals (29/42, red numbers in Figure 2f) showed a pattern that learned rapidly over the first 4 trials before plateauing across the remaining ones (Cluster 1, Figure 2c), whereas the others (blue individuals in Figure 2f) show a more gradual and slow learning pattern, slowly decreasing across trials (Cluster 2, Figure 2c). The slow vs. fast learning in this setting is highly dependent

on the initial performance of the individuals. The distinction between early and late trials can be confirmed by looking at the total within-cluster sums of squares using data from fewer trials (Figure 2d). This showed that running more trials was important to discriminate between the two groups. If the experiment had been stopped at, e.g. trial 6, it would be difficult to separate the individuals into two well-defined clusters, as we are able to do when looking at the full dataset.

An advantage of our analysis is that we can track individual learning patterns across trials. For example in Trials 1-4 there is a relatively even split in the numbers of individuals showing rapid learning (red, $n = 23$) and gradual learning (blue, $n = 19$) (Figure 2e). This proportion changes with the addition of extra trials (Figure 2f). Some individuals change their pattern across trials (e.g. individuals 32, 12 & 29, clustered as slow learners (blue) for Trials 1-4 but as fast learners (red) across Trials 1-12). Others continue with the same pattern throughout (e.g. individuals 5, 14 and 39 remain as gradual (blue) learners).

Sea Hero Quest

Here we observe a non-linear behaviour of the task scores across trials, as the trials are semi-independent of each other. This is why we opted to model this behaviour using b-splines. We do have few trials (only five), and could include the trial predictor as a categorical factor in our model, but we opted to use random b-splines per individual to showcase the use of these random effects. Interpretation of the b-spline estimated coefficients is now not as straightforward as in the previous examples. In fact, they introduce great flexibility at the cost of direct interpretability of the estimated coefficients. Notwithstanding, we can still look at the fitted curves and make inference the same way as before. Since this is a slightly more complex model, we experienced convergence issues when fitting it to the data when including a random b-spline with three degrees of freedom per individual. However, the estimates obtained were sensible and reproduced the behaviour in the data well at marginal and individual level. We performed the likelihood-ratio tests based on models fitted with random intercepts and slopes, which converged, allowing for the computation of log-likelihoods. The interaction between the nonlinear effects of trial and age were highly significant both for the linear predictors for the mean ($LR = 112.08$, d.f. = 9, $p < 0.0001$) and dispersion ($LR = 123.6$, d.f. = 9, $p < 0.0001$), suggesting that different age groups are associated with different nonlinear learning behaviour (see Figure 3a).

The output of the cluster analysis with four clusters clearly shows that age group dominates the clustering. This is noteworthy because age group plays no role within the clustering algorithm itself, only the individual-level random effects do, and still we see a clear separation by age group, with individuals 18-40 in cluster 1, 21-60 in cluster 2, 41-80 in cluster 3 and 61-80 in cluster 4 (see Figure 3b). An interesting feature of the cluster analysis (Figure 3c) is that there are many individuals that have a learning pattern not typical for their age group. For example, there are some individuals in the 21-40 age category (green) that show learning patterns more similar to the 18-20 age group (red). In addition, the 41-60 age category (blue) is split fairly evenly, with a group of individuals showing learning patterns more similar to a

younger age group (red and green) and a second group showing patterns more similar to the 61-80 age group (purple). Using this level of analysis may allow us to identify individuals with spatial learning issues early. This split in the 41-60 year old cohort suggests that the data is best divided into 3 rather than 4 clusters (Figure 3d).

Finally, we analysed each of the 4 experiments using linear models (see supplementary Table) to assess whether GLMMs would reveal anything extra at the overall (non-individual) level. In general, the results across both methods were comparable.

Discussion

Here we show that the use of GLMMs can be broadened, so that individual variation in a dataset is not just controlled for but used to provide a deeper analysis on individuals' performances. While we have applied our approach to discrete, continuous, censored and non-censored data, to demonstrate the possibility of using it across multiple datasets, the analysis may be suitable for datasets beyond these. Furthermore, while we have focused on learning among individuals and whether patterns of learning could be identified, a similar type of analysis may equally be applied to other psychological constructs and across other fields (e.g. examination of individual growth patterns in plants, effectiveness of a drug or therapeutic programmes on individual patients). In terms of our face-name associative learning task, we were able to report significant effects for condition (massed vs spaced) for the online experiment but not for the in-person experiment. More than this, we were able to examine individual learning patterns, and show where an individual may lie along the slope/intercept and whether they exhibit outlying behaviour. Further, individuals' performances can be compared across conditions as well as within conditions. For example, with the in-person experiment, individuals in the spaced condition showed more extreme patterns, with many displaying very good learning (high intercept and slope values) compared to those in the massed condition.

Although cluster analysis did not reveal any significant patterns in the face-name task, two distinct patterns of learning emerged with the virtual navigation NavWell task. One pattern showed an initial rapid period of learning, plateauing at strong performance that is sustained for the rest of the trials. The second cluster showed a slower and more gradual learning pattern that sees individuals learn across all the trials. This finding may support recent research showing that humans vary substantially in their spatial learning (Newcombe, 2018). Using a virtual college campus, Weisberg and Newcombe (2016) were able to classify individuals as integrators, non-integrators and imprecise navigators depending on whether individuals were able to form a cognitive map from different environments or not. Our results would support the idea of individual variation among individuals, showing both patterns of fast and slower learners. Although the water maze task is relatively simple, it is one of the most popular tools used to study learning and memory in both humans and non-human animal research (Thornberry et al., 2020; Vorhees & Williams, 2006), as well as being an important assay for different diseases and neurological conditions (Goodrich-Hunsaker et al.,

2010). Despite its popularity group results rather than individual patterns of learning are typically reported (e.g. Woolley et al., 2015 and our own work, Farina et al., 2015). As such, important information underlying the mechanism of spatial learning may be missed. The two patterns of water maze learning may be supported by different neural patterns. For example, it is known that the hippocampus supports fast learning, while other structures including the striatum and neocortex support a slower learning (Kumaran et al., 2016). However, learning is dynamic and this is reflected in individuals changing from one pattern to another across trials; again, this may be indicative of dynamic neural interactions between different learning systems (see for example Kosaki et al., 2015). Further investigation is warranted to examine the underlying biological and/or environmental factors that support the different patterns.

One of the interesting findings to emerge from the Sea Hero Quest data is that learning patterns show a good correspondence to the age group. As might be expected both the 18-20 and 21-40 year olds have similar learning patterns and show good learning across the five different levels of difficulty, especially when compared to the older groups. This is not unexpected as spatial cognition declines with age (Techentin et al., 2014), and the Irish sample fits readily to what has been observed internationally with the larger Sea Hero Quest data (see Coutrot et al., 2018). However, here we can examine individual performances in more detail and check to see how such individuals compare to their matched-aged cohort. We show that some individuals show patterns that more closely match a younger or older cohort, rather than their own. For example, the 41-60 age group is particularly interesting as some individuals show patterns that are more similar to a younger cohort (21-40), while many others show patterns more similar to the older group (61-80). Indeed our cluster analysis show three clusters may better account for the data than four. This age group may therefore represent an important transition period in terms of spatial learning. Further research is needed to examine whether this transition is similar for other cognitive processes. In addition, as spatial learning deficits may be an early indicator of dementia and Alzheimer's disease (Coughlan et al., 2018), being able to identify individuals in a younger age cohort that perform as if much older may help to provide an early warning and may be useful as a cognitive digital marker for future disease. However, we need to be cautious as the age groups chosen for our analysis were wide, spanning 20 years (e.g. 41-60) for the most parts (except for the 18-20 cohort); therefore, it is possible that those in their early 40s show similar patterns to the 21-40 age group. Similarly, those in their late 50s may show similar patterns to the 61-80 age group; further analyses need to be done to examine this. It is also important to note that the Sea Hero Quest data presents only four distinct trials, and therefore using splines to model the nonlinear behaviour may lead to overfitting. However, the objectives here were to describe these trends as best as possible through individual random effects, so that the subsequent cluster analyses could reveal learning patterns with respect to the participants' ages. If, in contrast, the objectives had been to predict responses for different trials, then a less flexible linear predictor could be more adequate.

In conclusion, we have demonstrated the flexibility of using GLMMs when using repeated measures. Examination of individual differences are important to identify outliers or simply those that perform well or not. Such an approach may be useful for educators or in a clinical setting to help identify those that might need further assistance or attention. Furthermore,

identifying clusters of individuals and learning patterns may allow further investigation of the underlying biological, environmental and other factors that may help explain why such patterns occur.

References

- Allen, M., Poggiali, D., Whitaker, K. et al. (2019). Raincloud plots: a multi-platform tool for robust data visualization. *Wellcome Open Res*, 4, 63.
- Barnhart, C.D., Yang, D., Lein, P.J. (2015). Using the Morris water maze to assess spatial learning and memory in weanling mice. *PLoS One*, 10(4):e0124521.
- Bates, D., Maechler, M., Bolker, B.M., Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67(1), 1-48.
- Bolker, B.M., Brooks, M.E., Clark, C.J., Geange, S.W., Poulsen, J.R., Stevens, M.H.M., White, J.-S.S. (2008). Generalized linear mixed models: a practical guide for ecology and evolution. *Trends in Ecology and Evolution*, 24,(3) 127-135.
- Bootsma, J.M., Hortobágyi, T., Rothwell, J.C., Caljouw, S.R. (2018). The Role of Task Difficulty in Learning a Visuomotor Skill. *Med Sci Sports Exerc*. 50(9), 1842-1849.
- Braithwaite, D.W., Leib, E.R., Siegler, R.S., McMullen, J. (2019). Individual differences in fraction arithmetic learning. *Cogn Psychol*, 112, 81-98.
- Brooks, M.E., Kristensen, K., van Benthem, K.J., Magnusson, A., Berg, C.W., Nielsen, A., Skaug, H.J., Maechler, M., Bolker, B.M. (2017). glmmTMB Balances Speed and Flexibility Among Packages for Zero-inflated Generalized Linear Mixed Modeling. *The R Journal*, 9(2), 378-400.
- Chiang, S., Haut, S.R., Ferastraoar, V., Rao, V.R., Baud, M.O., Theodore, W.H., Moss, R., Goldenholz, D.M. (2020). Individualizing the definition of seizure clusters based on temporal clustering analysis. *Epilepsy Res*. 163,106330.
- Chou, M.Y., Nishita, Y., Nakagawa, T., Tange, C., Tomida, M., Shimokata, H., Otsuka, R., Chen, L.K., Arai, H. (2019). Role of gait speed and grip strength in predicting 10-year cognitive decline among community-dwelling older people. *BMC Geriatr*. 19(1), 186.
- Cochrane, C., Ba, D., Klerman, E.B., St Hilaire, M.A. (2021). An ensemble mixed effects model of sleep loss and performance. *J Theor Biol*. 509, 110497.
- Commins, S., Duffin, J., Chaves, K., Leahy, D., Corcoran, K., Caffrey, M., Keenan, L., Finan, D., Thornberry, C. (2020). NavWell: A simplified virtual-reality platform for spatial navigation and memory experiments. *Behav Res Methods*, 52(3), 1189-1207.
- Coughlan, G., Laczó, J., Hort, J., Minihane, A.M., Hornberger, M. (2018). Spatial navigation deficits - overlooked cognitive marker for preclinical Alzheimer disease? *Nat Rev Neurol*. 14(8), 496-506.

Coutrot, A., Silva, R., Manley, E., de Cothi, W., Sami, S., Bohbot, V.D., Wiener, J.M., Hölscher, C., Dalton, R.C., Hornberger, M., Spiers, H.J. (2018). Global Determinants of Navigation Ability. *Curr Biol.* 28(17), 2861-2866.e4.

Demétrio, C.G.B., Hinde, J., Moral, R.A. (2014). Models for overdispersed data in entomology. In: Ferreira CP, Godoy WAC (Eds) Ecological modelling applied to entomology. Springer.

Farina, F.R., Burke, T., Coyle, D., Jeter, K., McGee, M., O'Connell, J., Taheny, D., Commins, S. (2015). Learning efficiency: The influence of cue salience during spatial navigation. *Behav Processes.* 116, 17-27.

Firth, J., Firth, J.A., Stubbs, B., Vancampfort, D., Schuch, F.B., Hallgren, M., Veronese, N., Yung, A.R., Sarris, J. (2018). Association Between Muscular Strength and Cognition in People With Major Depression or Bipolar Disorder and Healthy Controls. *JAMA Psychiatry.* 75(7), 740-746.

Goodrich-Hunsaker, N.J., Livingstone, S.A., Skelton, R.W., Hopkins, R.O. (2010). Spatial deficits in a virtual water maze in amnesic participants with hippocampal damage. *Hippocampus.* 20(4), 481-91.

Kliegl R, Wie P, Dambacher M, Yan M, Zhou X. (2011) Experimental Effects and Individual Differences in Linear Mixed Models: Estimating the Relationship between Spatial, Object, and Attraction Effects in Visual Attention. *Front Psychol.* 1:238.

Kosaki, Y., Poulter, S. L., Austen, J. M., & McGregor, A. (2015). Dorsolateral striatal lesions impair navigation based on landmark-goal vectors but facilitate spatial learning based on a 'cognitive map'. *Learning and Memory*, 22, 179-191.

Kumaran, D., Hassabis, D., McClelland, J.L. (2016) What learning systems do intelligent agents need? Complementary learning systems theory updated. *Trends in Cognitive Sciences.* 20(7):512–534.

Mattei, J., Bigornia, S.J., Sotos-Prieto, M., Scott, T., Gao, X., Tucker, K.L. (2019) The Mediterranean Diet and 2-Year Change in Cognitive Function by Status of Type 2 Diabetes and Glycemic Control. *Diabetes Care.* 42(8), 1372-1379.

McDade, E., Wang, G., Gordon, B.A., Hassenstab, J., Benzinger, T.L.S, et al. (2018). Longitudinal cognitive and biomarker changes in dominantly inherited Alzheimer disease. *Neurology.* 91(14), e1295-e1306.

Morris, R.G. (1981). Spatial localization does not require the presence of local cues. *Learning and Motivation*, 12, 239–260.

Newcombe, N.S. (2018) Individual variation in human navigation. *Current Biology*, 28, R952–R1008.

Paulus, M.P., Squeglia, L.M., Bagot, K., Jacobus, J., Kuplicki, R., Breslin, F.J., Bodurka, J., Morris, A.S., Thompson, W.K., Bartsch, H., Tapert, S.F. (2019). Screen media activity and brain structure in youth: Evidence for diverse structural correlation networks from the ABCD study. *Neuroimage*, 185, 140-153.

Quesque, F., Coutrot, A., Cox, S., de Souza, L.C. et al. (2020). Culture shapes our understanding of others' thoughts and emotions: An investigation across 12 countries, *PsyArXiv*, <https://psyarxiv.com/tg2ay/>

R Core Team (2021) R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.

Raboyeau, G., Marcotte, K., Adrover-Roig, D., Ansaldi, A.I. (2010) Brain activation and lexical learning: the impact of learning phase and word type. *Neuroimage*. 49(3), 2850-61.

Rigby, R.A., Stasinopoulos, D.M. (2005) Generalized additive models for location, scale and shape,(with discussion), *Appl. Statist*, 54(3), 507-554.

Seidler, R.D., et al. (2015) Individual predictors of sensorimotor adaptability. *Front Syst Neurosci*. 9, 100.

Simon, G. (2001) Choosing a first-line antidepressant: equal on average does not mean equal for everyone. *JAMA*. 286(23), 3003-4.

Song, R., Xu, H., Dintica, C.S., Pan, K.Y., Qi, X., Buchman, A.S., Bennett, D.A., Xu, W. (2020). Associations Between Cardiovascular Risk, Structural Brain Changes, and Cognitive Decline. *J Am Coll Cardiol*. 75(20), 2525-2534.

Spiers, H.J., Coutrot, A., Hornberger, M. (2021) Explaining World-Wide Variation in Navigation Ability from Millions of People: Citizen Science Project Sea Hero Quest. *Topics in Cognitive Science*, in press.

Stasinopoulos, D.M., Rigby, R.A., Mortan, N. (2018) gamlss.cens: Fitting an Interval Response Variable Using 'gamlss.family' Distributions. R package version 5.0-1. <https://CRAN.R-project.org/package=gamlss.cens>

Techentin, C., Voyer, D., & Voyer, S.D. (2014). Spatial Abilities and Aging: A Meta-Analysis. *Experimental Aging Research*, 40, 395–425.

Thornberry, C., Cimadevilla, J.M., Commins, S. (2021) Virtual Morris water maze: opportunities and challenges. *Rev Neurosci*. 32(8), 887–903

Vorhees, C. V., & Williams, M. T. (2006). Morris water maze: procedures for assessing spatial and related forms of learning and memory. *Nature Protocols*, 1(2), 848–858.

Weisberg, S.M., Newcombe, N.S. (2016). How do (some) people make a cognitive map? Routes, places, and working memory. *J Exp Psych: Learn Mem Cogn*. 42, 768–785.

Woolley, D.G., Mantini, D., Coxon, J.P., D'Hooze, R., Swinnen, S.P., Wenderoth, N. (2015). Virtual water maze learning in human increases functional connectivity between posterior hippocampus and dorsal caudate. *Hum Brain Mapp*. 36(4), 1265-77.

Figure Legends

Figure 1: (a) Observed (black lines) and predicted (blue curves) scores for individuals under massed vs. spaced learning conditions for the online experiment. The shaded areas represent the 95% confidence intervals for the true mean scores, based on the binomial GLMM (see Table 1). The curves from the massed and spaced conditions are statistically different (test for different slopes: $LR = 12.12$, $d.f. = 1$, $p = 0.0005$). **(b)** Predicted random intercepts and slopes for each individual in the study, who was either at a massed or spaced learning framework, as well as 85%, 90% and 95% ellipsoids based on the bivariate normal distribution, for the online experiment. Points falling outside of the 85% ellipsoids have their individual numbers indicated as potential outliers. Area of the ellipsis is calculated as $A = \pi ab$, where a and b are the major and minor axes, respectively. Eccentricity is calculated as $E = (1 - b^2/a^2)^{1/2}$. **(c)** same as (a), but now for the in-person experiment; **(d)** same as (b) but now for the in-person experiment.

Figure 2: (a) Observed (black lines) and predicted (blue curves) scores for individuals that took the NavWell virtual navigation task. The shaded area represents the 95% confidence intervals for the true mean scores, based on the right-censored gamma GLMM (see Table 1). **(b)** Predicted random intercepts and slopes for each individual in the study, as well as 85%, 90% and 95% ellipsoids based on the bivariate normal distribution. **(c)** The two-cluster solution obtained when carrying out hierarchical clustering analysis based on the Euclidean distance and Ward's method, using the predicted random intercepts and slopes per individual obtained from the right-censored gamma GLMM fitted to the full dataset. Blue curves are estimated from loess regression and are meant as a visual aid only. **(d)** Total within-cluster sums of squares (TWSS) for one- to seven-cluster solutions obtained when carrying out the cluster analysis to subsets of the data (trials 1 – 4, 1 – 6, 1 – 8, 1 – 10, and 1 – 12). **(e)** Dendrogram and colour scheme representing the two-cluster solution obtained after fitting the model using trials 1 – 4 and **(f)** 1 – 12, i.e. the full dataset.

Figure 3: (a) Predicted scores (solid curves) for individuals of different age groups from the SeaHero Quest experiment, and associated 95% confidence intervals for the true mean scores (shaded areas) based on the gamma GLMM (see Table 1). **(b)** Observed scores for each individual split by cluster, obtained from the four-cluster solution of the hierarchical cluster analysis carried out using the random effects from the model, using the Euclidean distance matrix and Ward's method. **(c)** Dendrogram representing the full hierarchical clustering analysis, with colour scheme reflecting age group belonging for each individual. **(d)** Total within-cluster sums of squares (TWSS) for one- to seven-cluster solutions obtained when carrying out the cluster analysis to the full dataset.

Figure 1

Figure 2

Figure 3