The Importance of Group Specification in Computational Modelling of Behaviour

Vincent Valton¹, Toby Wise² & Oliver J Robinson¹

¹ Neuroscience and Mental Health Group, Institute of Cognitive Neuroscience, University College London

² Max Planck UCL Centre for Computational Psychiatry and Ageing Research, University College London; Division of the Humanities and Social Sciences, California Institute of Technology

Introduction: Hierarchical model fitting has become commonplace in cognitive neuroscience. However, these techniques require us to formalise assumptions about the data-generating process at the group level, which may not be known. In computational psychiatry, we frequently encounter situations with multiple distinct groups of subjects, which makes this specification non-trivial. Specifically, we must choose whether to assume all subjects are drawn from a common population, or to model them as deriving from separate populations. These assumptions have profound implications for computational psychiatry, as they will undoubtedly affect the resulting inference (parameter recovery) and conflate or mask true group-level differences.

Methods: We used simulations to test these assumptions on synthetic multi-group behavioural data from commonly used multi-armed bandit tasks. We first demonstrate the influence of hierarchical vs non hierarchical approaches on parameter recovery. We then examine recovery of group differences in parameter values under different modelling approaches: (1) modelling groups under a common shared group-level prior (assuming all participants are generated from a common distribution, and are likely to share common characteristics); (2) modelling separate groups based on symptomatology or diagnostic labels, resulting in separate group-level priors. We evaluate the robustness of these approaches to variations in data quality and in prior specification.

Results: We confirm, firstly, that hierarchical approaches lead to better parameter recovery. Secondly, we argue that fitting groups separately (Model 2) provides the most robust inference across all conditions and perturbations. Thirdly, we confirm these assumptions have less of an impact with higher numbers of trials and subjects. Finally, we show that the advice to fit groups separately can fall down when extreme assumptions are made about group variance.

Conclusions: Our results suggest that when dealing with data from multiple clinical groups, researchers should use a hierarchical approach to parameter fitting, but account for hypothesised group differences in the model fitting procedure. In other words, under most normal circumstances, researchers should fit patient and control groups separately.

Introduction

Computational modelling approaches are becoming increasingly popular for the analysis of behavioral task data (Ahn, Haines, and Zhang 2017). For example, reinforcement learning models have been adopted to explain human and non-human learning behaviour (Dayan and Daw 2008); prospect theory models have been used to explain decision making biases (Schonberg, Fox, and Poldrack 2011); and drift diffusion models have been used to explain choice reaction times (Ratcliff et al. 2016). While the underlying structure of these models is quite different, the approach used to determine the explanatory power - or 'fit' - of such models is generally the same. Specifically, the 'free parameters' in each model (for example, the *learning rate* in a reinforcement-learning model; *risk aversion* in a prospect theory model; or *drift rate* in a drift diffusion model) are 'solved' for each individual in a process known as parameter fitting; and then various metrics are used to judge how well these parameters explain the actual data for a given individual or group. For instance, an 'information criterion' such as Bayesian Information Criterion (BIC) might be used to trade off goodness of fit with model complexity (Wilson and Collins 2019). Data might then be simulated for an individual using their optimised parameters, which is then compared to their real choices. If the model is good then the real and simulated data should be tightly correlated. An excellent summary of this approach can be found in (Daw 2011; Wilson and Collins 2019) and well-documented toolboxes now make it considerably easier for researchers to adopt this approach (Ahn, Haines, and Zhang 2017).

Such modelling approaches are valuable because they may ultimately enable researchers to bridge the gap between brain and behaviour (Huys, Maia, and Frank 2016). They can generate and test *mechanistic* predictions about how behaviour is generated which can, moreover, be constrained by 'biophysical plausibility'. That is to say, the computations used in the models can be restricted to those that could plausibly be generated by the firing and modulation of neurons. Reinforcement learning models, for instance, provided key predictions about prediction error signalling that was borne out by subsequent analysis of dopamine firing within the ventral striatum (Badre and Frank 2012).

An area in which such mechanistic understanding may prove especially valuable is psychiatry (Huys, Maia, and Frank 2016), where behavioural symptoms are currently the primary means of diagnosis *without* an understanding of their underlying causal mechanisms. Current first-line treatments, for instance, were developed through trial and error without an understanding of the underlying neurobiology and as a result, treatment selection is untargeted (recovery rates are less than 50% (Clark et al. 2018)) and development of novel interventions has stalled (Scannell and Bosley 2016).

A growing sub-discipline of 'computational psychiatry', therefore attempts to formalise differences in symptoms between healthy individuals and those with psychiatric symptoms in terms of the parameters of computational models (Huys, Maia, and Frank 2016; Robinson and Chase 2017; Adams, Huys, and Roiser 2016). The hope is that by formalising the computations by which

pathological symptoms emerge using computational models, we will better understand the underlying mechanisms of these symptoms; better understand the role of current treatments in modulating these mechanisms; and hence improve our ability to target and develop new treatments.

Specification of the problem

In the computational psychiatry approach we are not only interested in the free parameters of our models, but also how the specific parameters might differ across groups; for instance patients and controls. However, the precise nature of the sample is important because it has been shown that estimates of free parameters for individuals are greatly improved by leveraging group level information when fitting individual data (Daw 2011; Brown et al. 2020; Ahn, Haines, and Zhang 2017). Specifically, constraining individual parameter fits by group mean and distribution information minimises extreme and incorrect parameter assumptions and leads to better overall fits (Daw 2011)(**Figure 1**). In the first study of this paper we illustrate this point using simulations to show that in situations commonly encountered in psychiatric research - 1) low subject numbers due to difficult to recruit populations and 2) low trial numbers due to limited testing time - the use of a hierarchical approach results in considerably better parameter recovery.

Figure 1: illustration of the general hierarchical modelling approach, whereby group level information is used to inform individual parameter fits

However, this hierarchical approach reveals a conundrum. It requires the researcher to specify the population structure of the data. In a classic psychology experiment this may be relatively easy - all the participants are, for example, from a single university undergraduate course or online testing platform - so we might consider them emerging from the same, single, population. However, in a typical clinical study we are interested in how *different* populations of individuals (or the same individuals under different context, say on and off drug) perform on some measure of interest. Crucially, these are deliberately selected such that the sample as a whole has a bimodal distribution on symptom scores. So, do we therefore consider patients and controls as emerging from the same group - perhaps as extremes of single normal distribution - or do we consider them emerging from separate groups - perhaps as separate, overlapping, distributions (e.g. (Aylward et al. 2019;

Mkrtchian et al. 2017)))? The answer to this question might actually have a considerable impact on our final parameters. For instance, if we consider them emerging from a single distribution, then hierarchical approaches will serve to 'pull' all individual parameters closer together and hence obscure potential group differences. If we assume they come from different distributions, then parameters within each distribution will be pulled closer together, and potentially away from each other, and we may therefore overestimate group differences in these parameters.

The extent to which the hierarchical approach is robust to these assumptions has not, however, been comprehensively explored. Of note, this question is not the same as asking whether continuous variables (e.g. depression symptoms) correlate with model parameters. For this, Bayesian regression approaches have been developed (Moutoussis, Hopkins, and Dolan 2018; Boehm, Steingroever, and Wagenmakers 2018)(see discussion section for further comparison). Here, we address the question of how to model groups that are clearly distinguishable from one another, such as subjects with different diagnoses. Given the relative ease with which researchers can now adopt this hierarchical approach (Ahn, Haines, and Zhang 2017), a clearer delineation of the potential pitfalls is, we believe, timely. In the second study of this paper we therefore simulate the impact that assumptions about groups has on parameter recovery, demonstrating, again, that some common assumptions may be particularly problematic in situations commonly encountered in psychiatric research.

General Methods

Reinforcement learning model

Reinforcement learning models can describe how associations between stimuli and outcomes are learned. For instance, RL agents can learn that some stimuli lead to bad, or punishing outcomes, while others lead to good or rewarding outcomes. There are many papers which provide a thorough discussion of RL, its implementation in the human brain, and how it relates to psychiatric symptoms (Wilson and Collins 2019; Daw 2011; Robinson and Chase 2017; Adams, Huys, and Roiser 2016; Huys, Maia, and Frank 2016) so we do not cover it in detail here. However, throughout this paper we generate and recover data from a simple reinforcement learning (RL) model which generates action probabilities from outcomes using two free parameters (learning rate and temperature). In mathematical terms:

$$Value_{\{s_t\}(a_t)} \leftarrow Value_{\{s_t\}(a_t)} + LearningRate \cdot (outcome_{\{t\}} - Value_{\{s_t\}(a_t)}) \quad (1)$$

$$ActionProbability_{\{s_t\}(a_t)} = \left[\frac{\exp\left(-\frac{Value_{\{s_t\}(a_t)}}{Temperature}\right)}{\sum_{b=1}^n \exp\left(-\frac{Value_{\{s_t\}(a_t)}}{Temperature}\right)} \right] \quad (2)$$

Where *outcomes* are determined by the task (1 for a reward, 0 for no outcome, -1 for a punishment), *Values* are initialised at the mid point (e.g. 0 if both rewards and punishments are present) and change on each trial (t) depending on outcomes and the *Learning Rate* which describes how quickly information is acquired. These are all used to determine the probability of choosing a given option (*Action Probability*) given a *Temperature* parameter, which describes how stochastic (or noisy) choices made on the basis of this *Value* information are.

Groups

In our simulations we assume two labelled and meaningful categorical groups; for instance patients and controls. It should be noted that this is probably a best case scenario; it is possible that psychiatric groupings do not represent ‘natural kinds’. That is to say psychiatric diagnosis may not actually reflect true, delinatable groups. However, in our simulations, we systematically vary the size of group differences such that we can see the impact of group assumptions on no (d=0), to very small (d=0.05) to extremely large (d=5) group differences. Psychiatric group differences in parameter estimates are likely to be more towards the small end (d<0.2).

Multi-armed bandit task

We tested the model on two tasks, a two or four armed bandit task. On each trial, participants are asked to select one of either two or four bandits and are then provided (for just the selected bandit) with one of: 1) no feedback (0), 2) positive feedback (+1), 3) negative feedback (-1), or 4) both

positive and negative feedback. The probabilities of these outcomes fluctuated independently and rapidly across bandits, such that the bandit that was most rewarding changed over time (**Figure 2**)

Figure 2 Illustration of the task used for simulations. (a) In each trial of their task, subjects selected one of the four bandits and received a reward based on its current mean pay-off perturbed by noise. (b) The expected value of each bandit changed continuously over time.

Such tasks have previously been used to assess differences between clinical groups (Aylward et al. 2019). We chose to use roving probabilities as this allows us to estimate learning rate and temperature more accurately than conventional stationary RL environments. Specifically, traditional stationary RL environments (e.g. with fixed 80/20 probability of reward) result in fewer trials being usable to estimate learning rates, as participants will tend to stay with the best option once they've identified it. Using a multi-arm bandit with roving probabilities allows us to ensure maximal parameter recovery for both exploration and learning rate parameters, and to avoid introducing biases in the parameter recovery due to asymptotic performance at the task (Seymour et al. 2012).

Parameter fitting

We compare a number of different approaches to fitting parameters. These approaches have all previously been used in papers in this field and all have advantages and disadvantages.

Maximum Likelihood Estimation (MLE)

MLE is a computationally simple method for parameter estimation that does not incorporate any information about prior distributions on model parameters (for more details see (Daw 2011)). While fast and often accurate for simpler models, MLE has two major disadvantages: first, it is not best suited for more complex models, for example those with a hierarchical structure. Second, it only provides point estimates of parameters without any information about the uncertainty around these estimates. While we note that maximum-a-posteriori (MAP) estimation is also available as an alternative that allows prior specification (MLE is in fact a special case of MAP), this also only provides point estimates and does not perform well with complex models.

Markov chain Monte Carlo (MCMC) Sampling

MCMC sampling is a method that aims to approximate the posterior distribution of model parameter values (Kruschke 2014). At their core, sampling algorithms evaluate the likelihood of the data given a particular set of parameter values, which are iteratively sampled from the parameter space. Sampling is biased towards sampling effective parameter values more frequently, and as a result the sampled parameter values provide an accurate approximation of the true posterior distribution. Sampling is able to accurately estimate posteriors for complex models, and provides the full posterior distribution over their values. Importantly, being a Bayesian method, MCMC sampling requires priors to be placed on parameter estimates. However, this approach is computationally highly intensive.

Hierarchical Markov chain Monte Carlo (hMCMC)

In this paper when we specify the full hierarchical approach (i.e. taking into consideration the groups) we do so using the MCMC approach. We therefore use hMCMC to refer to hierarchical models estimated using MCMC sampling. For more information about this approach see (Kruschke 2014; Gelman et al. 2013). Again, however, this approach has the disadvantage of being highly computationally intensive.

Hierarchical Variational Bayesian Inference (VB)

Variational inference (VB) provides a method to approximate posterior distributions of model parameter values without using sampling (Penny, Kiebel, and Friston 2003), and thus represents a faster, but potentially less accurate, method of estimating posteriors. Here we use this approach for especially computationally intensive hierarchical simulations after having demonstrated (in study I) that hierarchical VB does a good job of approximating parameters from the hMCMC approach.

Simulations

We ran two separate simulations denoted as studies I and II. In study I, we explored the impact of the different parameter fitting procedures for different numbers of participants and trials acquired.

In study II, we explored the effect of group priors in hierarchical approaches, again as a function of participants and number of trials.

Study I: Rationale for Hierarchical approaches to parameter fitting

Introduction

As highlighted in the introduction, hierarchical fitting approaches aim to improve parameter estimation by using group-level parameter distributions to inform those of the individual subjects' (Gelman et al. 2013; Daw 2011). In this section, we demonstrate the benefits of hierarchical estimation under different, commonly experienced, contexts. For this illustration, we focus on a relatively simple use case: estimating RL parameter values on a binary choice task (i.e. a two armed bandit: **Figure 2** but with only two options). Our aim here is to show how, even with tasks and models that are low in complexity, hierarchical fitting can provide substantial improvements in parameter estimation over other methods. We compare these approaches on their ability to accurately recover parameter values used to generate simulated data across a range of subject and trial numbers.

Study I specific methods

Task

The task used for these simulations was a simple two-armed bandit task, with two options offering reward probabilities of 0.2 or 0.8 which reversed every 40 trials. Similar tasks are commonly used in computational psychiatry and are often described well by standard reinforcement learning models.

Data simulation

The reinforcement learning model comprised learning rate and softmax temperature as two free parameters (**see Equations 1 and 2 above**). Simulations were run using 9 combinations of parameter values for the group mean, comprised of learning rate values of [0.3, 0.5, 0.7] and temperature values of [1, 3, 5]. Group standard deviations were 0.2 and 0.4 respectively. We repeated these simulations for group sizes of $n=15$ and $n=200$, to estimate recoverability across small and large group sizes, and with $t=40$, $t=200$, and $t=300$ trials to investigate how trial numbers affect recoverability. While some of these numbers are small, we note that it is not uncommon to have small sample sizes and trial numbers in clinical studies due to difficulties in recruitment, and time pressure once patients are recruited.

Model fitting

All models were fit using PyMC3 (<https://docs.pymc.io/>), a probabilistic programming package in Python that implements MCMC sampling. MLE was also performed using PyMC3, based on optimisation routines implemented in Scipy (<https://www.scipy.org/>). For MCMC estimation, we used two chains of 3000 samples, while for variational inference we used 40000 iterations of the advanced differentiation variational inference algorithm.

For non-MLE approaches, we used wide priors, with means of 0.5 and 3 for learning rate and temperature respectively, and standard distributions of 100 for both. For hierarchical models, these priors were placed on the group-level mean, represented by a gaussian distribution with variance given as a uniform distribution with bounds at 0 and 100, while for non-hierarchical these were placed on the individual-level parameters directly.

To evaluate parameter recovery we compared the recovered parameter estimates to the true values used to simulate the data using mean squared error (MSE). MSE was averaged across different combinations of group mean parameter values. Greater MSE means worse parameter recovery.

Study I results

Result of parameter recovery analyses are shown in **Figure 3**. Across both temperature and learning rate in the model, and irrespective of the number of subjects and trials, hMCMC provided the most accurate parameter estimates, followed by hierarchical variational Bayesian inference and non-hierarchical MCMC. Maximum likelihood estimates were substantially less accurate, and this was accentuated with small trial numbers. Notably, hierarchical MCMC was able to recover parameters accurately even with very few subjects and very few trials, a situation that is frequently encountered in computational psychiatry.

Figure 3: Parameter recovery error (Mean Square Error MSE) is considerably greater in non hierarchical (MLE) approaches than fully hierarchical (hMCMC) approaches. This problem is especially bad when there are fewer trials (i.e. when data collection time is limited) and when there are smaller numbers of participants (i.e. when patients are hard to recruit).

Study I Summary

Together, these results demonstrate the power of hierarchical MCMC when estimating model parameters from behavioural in learning tasks. Even in situations with limited data, parameter estimates were still reasonably accurate using this approach. Both hierarchical variational Bayesian inference and non-hierarchical MCMC also provided accurate parameter estimates, however MLE estimates were less accurate, particularly when few trials were available. Importantly, the model and task here are among the least complex used in the field; the task only requires a single value estimate to be learned, and we only need to model binary choices, using a model that only has two free parameters. Typically, we will be using more complicated models and tasks in real studies, and so the fact that MLE provides poor parameter estimates even in this simple setting suggests that it may fail to a greater extent in more (typical) complex use cases.

Overall, study I indicates that hierarchical estimation should be the preferred method of parameter estimation whenever feasible to use, and especially when participant and/or trial numbers are limited. In situations where MCMC is infeasible, variational inference provides a useful and accurate approximation.

Study II: The importance of model specifications in hierarchically fitted data

Introduction

As highlighted in the introduction, however, the hierarchical approach favoured by study I raises some empirical problems when it comes to psychiatry. Critically, we do not know whether patients and controls (classified using diagnostic criteria based on symptoms) are actually coming from two distinct populations (patients and controls), or are actually coming from a common distribution with inter-individual variability such that patients and controls all lie within a spectrum/continuum in symptom and parameter space.

Such assumptions have profound implications for hierarchical modelling paradigms, as one needs to incorporate all known underlying assumptions within the model. Practically speaking, this issue translates into a problem of model definition. That is, shall we:

- **Model 1:** model separate groups under common shared group-level priors (making the assumptions that participants are generated from a common distribution, and are likely to share common characteristics).
- **Model 2:** model separate groups based on symptomatology (i.e. based on DSM or ICD-10 clinical diagnosis) such that they belong to separate groups in the parameter-model space (i.e. end-up having different group-level priors).
- **Model 3:** assume that different groups should be modelled separately, but that these groups all come from an underlying generative distribution such that groups are more likely to be similar to each other (i.e. modelling the assumption of a continuum between groups). This would be most similar to an RDoC approach, where symptoms are believed to lie on a continuum from healthy to pathological. That is we enforce similarity at the group-level rather than at the individual-level across group.

The assumptions of each of these models (**Figure 4**) has profound implications for computational psychiatry and may very well affect the inference made (e.g. how well we are able to recover individual parameters) and conflate or diminish group level differences (e.g. how often we observe significant differences between groups on model and/or parameter space).

Figure 4. Illustration of different model specification approaches for data from two psychiatric groups.

In this section we use simulations to test these assumptions on synthetic datasets, in an attempt to answer the following primary questions:

Question 1: What are the best assumptions with respect to group-level priors, in order to best recover the within-group level parameters?

Question 2: How does our best hierarchical model assumptions fare in edge cases such as with small sample sizes ($n=15$), or few available trials per participant ($t=40$).

Question 3: Can we identify situations where answers to the above questions break down?

Study II specific methods

Synthetic data

We generated data using a typical RL model for two separate groups, with varying degrees of overlap between parameters across groups. To study the effect of different model specifications on model fit and parameter recovery we opted to start with a sufficiently large number of subjects per group ($n=50$), and number of trials ($t=200$). This ensured that differences between recovery of model parameters were due to model specifications alone, and not due to insufficient data for parameter recovery.

We parameterized the synthetic dataset such that learning rate and temperature always fell between zero and one. Beta distributions were used to ensure these constraints, and reparameterized with different mode and concentration parameters. We then generated 36 synthetic datasets by modifying the two groups mode between the ranges of $[0.01, 0.1, 0.2, 0.3, 0.4, 0.5]$ for group one, and mode of range $[0.99, 0.9, 0.8, 0.7, 0.6, 0.5]$ for group two respectively, with concentration parameters varying in the range of $[2, 2.5, 3, 5, 10, 30]$. This range of values allowed us to simulate groups with complete overlap in parameter space, with varying degrees of similarity within each group (see **figure 5** below). Using these datasets, we could then compare models of

different specifications, and study their failure points under different hierarchical assumptions or for different effect sizes. For every synthetic dataset generated, we ensured that the mean and standard deviation of the generated parameters matched this of the data-generating distribution (that is, to generate 50 participants learning rates with mean 0.5 and standard deviation of 0.2, we ensured the sample of random 50 learning rates was within 0.02 absolute error of the theoretical data-generating distribution). This was done so as to ensure that any bias observed during the recovery of group-level parameters (i.e. group mean and standard deviation) were due to the model specifications, and not due to random unrepresentative samples from the data generating distribution. In the main paper we simplify the data generating process such that all data had different learning rates, but shared the same temperature parameter (constant fixed to 0.8 -- see **Figure 5** below). For the effect of varying temperature or both see the supplement.

Figure 5: Distribution of learning rates for the 36 synthetic datasets in which only the learning rate differs between groups and temperature is fixed to 0.8 and is known (i.e. isn't estimated during parameter fitting). See supplement for more complex data sets

Varying data quality: number of subject per group and number of trial per subject

Finally, in order to assess the robustness of each model specification to varying levels of data quality we varied number of subjects and trials (Baker et al. 2019). Due to combinatorial issues, we opted to test only two group sizes, and two sets of trial size, which we agreed verged on the limits of acceptable, feasible and normative data quality. We assumed 50 subjects per group with 200 trials per subject as our optimal (High) data quality and allowed for a lower group size of merely 15

subjects per group (as it is not uncommon to see such sample sizes in costly experimental settings such as fMRI studies (T. Wise et al. 2017; Nord et al. 2017)). We opted for 40 trials as our lowest use-case for low data quality as it isn't unusual to see RL studies with no more than 2-4 pairs of stimuli with only 10-15 trials per pair in a stationary environment. In our datasets we only model one set of 4 stimuli for 40 trials, but the probability of reward and punishment is non-stationary, therefore allowing much better recovery than with multiple sets of stationary environments. We selected a sequence of 40 consecutive trials from the 200 trials data which allowed for highest recovery of parameters (ρ approx 0.8), therefore the analysis should provide an upper bound as to what each model is able to achieve by each method when using sub-optimal data.

All models were rigorously tested for quality checks on uniformly distributed synthetic data to ensure that all three model specifications could adequately recover parameters. Additionally for all models group-level and population-level priors were calibrated such that group-level priors were comparable across models (Gabry et al. 2019), and to ensure that any observed difference between models were due to the specification of hierarchy, but not due to prior-level choices.

Model fitting

Models were fit using a combination of hierarchical MCMC and VB implemented with the Matlab (Matlab 2012) interface to Stan (Gelman et al. 2013). Samplers used 2 chains of 3000 samples, with the first 1000 of each used for warm up and discarded. Group-level effects were specified as beta distributions (Learning rate $\sim$ Beta($N(1,5)$, $N(1,5)$): different hyper-parameterizations discussed in Question III).

Study II Results

One of the most important metrics we want to recover from our model fitting is the true difference between groups. To do so, we need to accurately recover the group-level summary statistics such that we can adequately compare groups means. Ideally this would be done using fully Bayesian approaches (Kruschke 2014) (see (Ahn, Haines, and Zhang 2017) for a tutorial on implementing some of these approaches), but in practice this is often pragmatically done using standard general linear models (e.g. t-tests, anovas, etc.). One way to visualise parameter recovery for both groups using a single metric is to therefore use the effect size of the difference between groups, as it naturally combines both group means and standard deviation, which themselves rely on the correct estimation of the subject-level parameters within each group. To compare each model specification, we therefore calculate the true effect-size (Cohen's d) between the two groups from the synthetic datasets, and then calculate the recovered effect-size between groups using each model specification.

We present here these results for the 36 datasets plotted in **Figure 5** (for further details of performance on other datasets, including typically underperforming results from model 3, see Supplement).

Question I: What are the best group priors to recover true effects?

We examined the variability of recovery using different models. Due to the randomness inherent in the generation of parameter values for simulations, it is possible that we might inadvertently produce datasets with particular characteristics that make them difficult for certain fitting approaches. As a result, we sought to repeat our analyses 1000 times for all 36 datasets, all with the same group mean and variance in the parameters of interest in order to rule out the possibility that our effects are due to peculiarities in the dataset. This also allows us to approximate the false positive and false negative rates in our analyses, as we can infer the number of examples in which we identify a difference that does not exist, or fail to identify one that does. Given the computational complexity of MCMC sampling, it was not feasible to repeat our analyses using this approach (total of 36,000 datasets). Instead, we used variational Bayesian inference to approximate the posteriors of our parameters. As shown in **Figure 1**, variational inference is capable of providing estimates that approach the accuracy of MCMC, albeit with substantially less computational demand.

Results of these analyses (**Figure 6**) show that fitting groups separately (model 2 in **Figure 4**) provides better estimates of group-level differences than fitting groups together (model 1 in **Figure 4**). Importantly, these results are not simply due to peculiarities in the datasets being used. For example, the confidence intervals for Model 2 all crossed equality, indicating that this approach does not result in any significant misestimation of the effect size. In contrast, almost all confidence intervals for Model 1 were robustly below the line of equality, indicating that this approach significantly underestimates the effect size.

Models

***Figure 6: Recovered effect-sizes** demonstrate that modelling all patients and controls under a single distribution results in substantial underestimation of group differences, whereas modelling the groups separately only slightly overestimates group differences (in all cases the variance overlaps perfect recovery). This data is generated using variational Bayes, see supplement for more complex data sets and models.*

In sum, the recovered effect sizes suggest that using a single prior over all groups (model 1) always underestimates the true effect-size between groups. This is because the group-level prior pulls the estimation of parameters for both groups to be more similar to one another. Modelling groups separately provides a much closer (albeit slightly overestimated) recovery of true effect sizes.

Question II: How does data quality influence effect size recovery?

We were next interested in determining how data quality impacted parameter recovery, as it could be that some model specifications are particularly prone to error when data quality deteriorates. We focused here on the number of trials and the number of subjects, two factors that are often limited in psychiatric research.

Model 1 was particularly inadequate in the poor data-quality cases (low number of trials), in which it underestimated the true effect size by as much as [-63:-65]% (depending on sample size; see **Figure 7**, bottom two rows). In more optimal data quality settings (high number of trials: **Figure 7** upper two rows), Model 1 still underestimated the true effect size, but by a much more acceptable margin [-18:-29]% (see Figure 6, upper two rows).

Conversely, using a different group-level prior for each group (model 2) allows 'pooling' solely for the individuals belonging to the group being estimated. As a result model 2 tends to most accurately recover the effect-size between the two groups (**Figure 7**). Similarly, data quality appears to degrade the recovery of the effect size (particularly so on data with a low number of trials). While model 2 appear to overestimate effect size by [+8:+10]% in high data quality settings (**Figure 7** upper two rows), it was not as catastrophically affected as model 1 in low data quality settings (**Figure 7** bottom two rows - overestimated effect size range [+28:+29%]).

Models

Data quality		# Participants	
# Trials	T=40	N=15	N=50
	T=200	Poor	Med.
		Med.	High

Data quality		# Participants	
# Trials	T=40	N=15	N=50
	T=200	Poor	Med.
		Med.	High

Data quality		# Participants	
# Trials	T=40	N=15	N=50
	T=200	Poor	Med.
		Med.	High

Data quality		# Participants	
# Trials	T=40	N=15	N=50
	T=200	Poor	Med.
		Med.	High

Figure 7: Recovered effect-sizes depend on data quality, with substantial underestimation of group differences when only a single group is assumed. In all cases recovery is better when the groups are modelled separately. See supplement for more complex data sets and models.

In summary, modelling groups separately is always preferred, but problems with modelling groups under a single prior are exacerbated by poor data quality, especially when the numbers of trials collected is low.

Question III: Can we identify situations where our advice breaks down?

It is also important to attempt to identify limits to our recommendations. To this end, the choice of hyperparameters for the group-level priors (which affect the assumptions about what we know about the group-level mean and the group-level standard deviation) may change the model 2's ability to accurately recover the effect size. For example, we might assume that the group-level mean is totally unknown (uniform prior over the group-level mean learning rate), but that the group-level standard deviation is known (we assume that subjects within a group are more similar to one another). Or conversely, we might assume that the group-level mean is known, but the group-level standard deviation is unknown. Both parametrizations result in a uniform group-level prior, but the implications for the model are vastly different. The first set of assumptions (unknown mean, known variance) will result in large levels of 'pooling' of participants parameters within their respective group, while the latter set of assumptions (known mean, unknown variance) results in little if any group-level 'partial-pooling'. It is therefore imperative that a reasonable set of hyperparameters is chosen when specifying the model. Of note, this issue would be completely missed when doing typical parameter recovery on a synthetic dataset where the range of parameter values span the whole range of possible parameter values and it is typically done on a single group.

Below we present the same data, recovered using a different hyper-parametrization such that the group-level mean is known, and group-level variance is unknown (i.e. Learning rate $\sim \text{Beta}(N(1,0.05), N(1,0.05))$ rather than Learning rate $\sim \text{Beta}(N(1,5), N(1,5))$ specified above), resulting in little if any group-level partial pooling. We show that even in this worst case scenario, model 2 underestimates the effect size but still less so than model 1 (**Figure 8**). Similar results were also found when varying the data quality using this hyper-parameterization, namely that model 2 still provided the most accurate model specification to recover effect sizes (**Figure 9**). It is worth mentioning however, that should researchers use hyper-parameterization that *enforces* large levels of partial-pooling (i.e. small group-level variance e.g. learning rate $\sim \text{Beta}(N(1,55), N(1, 55))$) model 2 can substantially overestimate the true effect-size (**Figure 10**). It is therefore imperative that researchers use appropriate and adequate hyperparameters for group-level shrinkage .

Models

Figure 8: Effect of hyper parameterization (assuming extremely wide group variance) on effect-sizes recovered using variational Bayes. As in figure 6, we demonstrate that modelling all patients and controls under a single distribution results in substantial underestimation of group differences, whereas modelling the groups separately only slightly overestimates group differences (in all cases this overlaps perfect recovery). See supplement for more complex data sets and models

Models

Data quality		# Participants	
# Trials		N=15	N=50
	T=40	Poor	Med.
	T=200	Med.	High

Data quality		# Participants	
# Trials		N=15	N=50
	T=40	Poor	Med.
	T=200	Med.	High

Data quality		# Participants	
# Trials		N=15	N=50
	T=40	Poor	Med.
	T=200	Med.	High

Data quality		# Participants	
# Trials		N=15	N=50
	T=40	Poor	Med.
	T=200	Med.	High

Figure 9: Effect of hyper-parameterization (assuming extremely wide group variance) and data quality on effect-sizes recovered using sampling. See supplement for more complex data sets and models. In no cases is model one preferred.

Figure 10: Effect of inadequate hyper-parameterization (assuming extremely small group variance) results in substantial overestimation of effect-sizes recovered using sampling for model 2. See supplement for more complex data sets and models.

In summary, modelling groups separately is usually preferred, however this does depend on the specification of hyperparameters. When some unlikely assumptions are made (e.g. that individuals within a group do not vary very much) then the recommendation may fall down.

Study II summary

In Study II, we sought to identify the most appropriate way to model data from multiple groups. Our results indicate that assuming all subjects are drawn from the same group (Model 1) leads to substantial underestimation of differences between groups due to shrinkage encouraged by the modelling approach, whereby all subjects are drawn towards a common group mean. In contrast, Model 2, which assumed subjects were drawn from separate groups, was able to recover the effect size of the simulated group difference accurately. Using variational inference, we were able to repeat this analysis across multiple simulated datasets, indicating that this effect was robust across many permutations of the data. The effect was most prominent with low numbers of trials, but persisted even with 200 trials. The recommendation breaks down if extreme assumptions are made about within-group variances, but generally holds under normal assumptions.

Overall, study II indicates that modelling groups separately generally provides the most accurate parameter estimates.

Discussion

Choosing an appropriate parameter estimation strategy is critical for ensuring accurate parameter estimates from computational models, and thus ensuring correct inference regarding computational dysfunction in psychiatric disorders. Here, we demonstrate that hierarchical models representing diagnostic groups *separately* provide the most accurate parameter estimates, and are able to most closely recapture group differences in parameter values. In contrast, models assuming that all individuals are drawn from a single group tend to (substantially) underestimate differences between groups.

Prior work has demonstrated the benefits of hierarchical models in other areas (Daw 2011; Wilson and Collins 2019; Brown et al. 2020), highlighting their ability to constrain parameter estimates such that they are less likely to be misestimated. In the present paper, we add to this by illustrating the benefits of hierarchical modelling (especially MCMC sampling approaches (Gelman et al. 2013)) to properly estimate these models in the context of RL tasks of the kind commonly used in computational psychiatry. Even with a simple task and model, we found that MLE, an estimation technique commonly used in the field, results in far less accurate parameter estimates than hierarchical modelling approaches.

Having shown that hierarchical models provide the most accurate parameter estimation, we next addressed the question of how to design models when we have multiple groups. This is a common problem in computational psychiatry, where we will often be comparing parameter values between diagnostic groups or treatment groups. However, it has remained unclear how best to capture this design in a hierarchical model. Intuitively, a model that assumes subjects are drawn from one group does not capture the true structure of the data, and the shrinkage caused by the hierarchical approach could lead to group parameter estimates being estimated as more similar to one another. However, modelling groups separately, while more reflective of the true data structure, has the potential disadvantage of artificially creating group differences when none exist. Our results suggest that concerns with the single group approach do emerge in reality, resulting in an underestimation of group differences. Modelling each group separately however results in far more accurate parameter estimates, with only minor exaggerations in estimated effect sizes (under most normal situations). Thus, when specifying hierarchical models for two separate diagnostic (or treatment) groups, at least in the case of a relatively simple 2 parameter reinforcement learning model, it is better to specify the groups under separate priors.

These results of course depended on data quality, most notably the number of trials. When tested with 40 trials, non-hierarchical approaches substantially increase recovery error. However, even when using a hierarchical approach, the model assuming all participants emerge from a single group dramatically underestimates differences between groups, while the approach modelling

groups separately was still able to provide accurate estimates of the group difference in this situation. While it is always recommended to use as many trials as possible, this may not always be feasible in clinical studies (where patients often have limited time to perform multiple tasks assessments). Our results emphasise the importance of modelling groups appropriately in these situations.

Taken together, our results have important implications for studies using computational modelling in situations with multiple groups of subjects. First, MLE provides far poorer parameter estimates when fitting computational models to behavioural data than hierarchical MCMC, and so should be avoided (as has been noted before (Brown et al. 2020)). While we acknowledge that in large samples this can be computationally difficult, we additionally showed that variational inference, a less demanding approximate inference method, can also provide accurate estimates for hierarchical models in these situations (Penny, Kiebel, and Friston 2003; Toby Wise and Dolan 2020). Second, when modelling multiple groups, these groups should be ideally modelled using separate group-level effects. That is, the hierarchical model should assume that subjects are drawn from their respective group, rather than from a single group.

We hope that these recommendations will be applicable across a range of studies within computational psychiatry, however we do urge researchers to be careful when performing any such analysis. In particular, we strongly recommend performing similar simulation-based checks to those reported here to ensure that the modelling approach selected can indeed recover hypothesised group differences in parameter values for the particular models used in a given study. The results reported here are based on relatively simple models, with at most two free parameters that are not highly correlated with one another. With more complex models, it is possible that different model structures may be more or less appropriate, and we advise that this is checked using simulation before fitting models to real data.

Although the issue of proper model specification in computational psychiatry has remained relatively under-researched, there are nonetheless important contributions that should be acknowledged. In particular, it has been previously suggested that the optimal way to investigate symptom-parameter relationships is to build them into the model (Moutoussis, Hopkins, and Dolan 2018; Boehm, Steingroever, and Wagenmakers 2018). This approach generally represents group-level parameter values as being offset from a single mean to an extent dependent on a psychiatric covariate e.g. a continuous psychiatric symptom covariate). However, this could be extended to the dichotomous case by coding the groups as, for example 0 and 1. We argue that this is largely equivalent to Model 2 which assumes separate priors for different diagnostic groups and, while this approach may be successful, and avoid shrinkage-related issues of under- or over-estimation of effect sizes, it can present difficulties. First, when using more than two groups, group identity will need to be dummy coded, which raises the issue of choosing the optimal dummy coding scheme. Second, this approach typically assumes that each group has identical variance. In contrast, our approach allows different groups to have different variances. Third, when multiple covariates are included, as is often the case in psychiatric research where potential confounds must be dealt with, and we wish to model these effects across multiple parameters of interest, model

complexity will increase substantially when this is built into the same model. As a result, parameters will inevitably be less identifiable and may trade off against one another. Separating the estimation of model parameters from the estimation of parameter-symptom relationships avoids this problem. Moreover, the two group approach from model 2 in this paper can be easily implemented using currently available toolboxes (Ahn, Haines, and Zhang 2017) simply by fitting the model to the data from each group separately (see e.g. (Aylward et al. 2019)).

While these results provide concrete suggestions for model fitting in this area, there are some caveats that should be considered. For example, hyperp parameters (those governing group-level parameter estimates) must be chosen carefully. While these should overall have a relatively limited effect with sufficient data, we show that they can nonetheless influence posterior estimates. Modelling groups separately tends to produce accurate parameter estimates except in the extreme scenario where one assumes extremely low group variance, in which case modelling groups separately leads to substantial overestimation of group differences. In practical terms it is unlikely researchers would choose such extreme hyperparameters, but it is valuable to know where our recommendation falls down. It is also important to note that our results are based purely on simulations. While these allow us to accurately evaluate parameter recovery, they are not entirely reflective of reality, where we are faced with additional sources of noise in behaviour due to attentional and motivational factors (i.e. measurement error). Another caveat worth mentioning, is that in order to compare extracted parameters using comparable methodology we estimated the group mean and group standard deviation from individual-level parameters for both model 1 and model 2. Ideally, for model 2 we would use the group level posteriors for inference (rather than individual subjects) as it would capture both individual-level and group-level uncertainty, but since these priors were unavailable with model 1 specification, we opted to use a similar methodology for model 2 so as to be comparable across models. Finally, it is worth noting that study II uses a combination of hierarchical MCMC and VB to provide inferences. This is due to the computational intensity of some of our simulations which rendered full MCMC approaches computationally intractable. Future work with better technology may seek to confirm predictions from VB using MCMC.

Our results raise important questions for future research. For example, while the approach of modelling groups separately provides the most accurate parameter estimates, this relies on the availability of known group labels. It may be possible to combine unsupervised clustering methods into the parameter estimation approach to allow the automatic identification of groups within the data, obviating the need for diagnostic groups to be manually assigned. In addition to potentially providing more accurate parameter estimates, this could permit the delineation of patient clusters based purely on computational modelling of behaviour.

Summary

Our results highlight the importance of assuming the correct hierarchical structure when fitting computational models to data from multiple groups. Based on our results, models should be fit hierarchically, with different patient (or intervention) groups fit separately. We hope that this

guidance will facilitate the identification of computational dysfunction in psychiatric conditions, and provide a starting point for further work on optimal approaches to modelling these populations.

References

- Adams, Rick A., Quentin J. M. Huys, and Jonathan P. Roiser. 2016. "Computational Psychiatry: Towards a Mathematically Informed Understanding of Mental Illness." *Journal of Neurology, Neurosurgery, and Psychiatry* 87 (1): 53–63.
- Ahn, Woo-Young, Nathaniel Haines, and Lei Zhang. 2017. "Revealing Neurocomputational Mechanisms of Reinforcement Learning and Decision-Making With the hBayesDM Package." *Computational Psychiatry (Cambridge, Mass.)* 1 (October): 24–57.
- Aylward, Jessica, Vincent Valton, Woo-Young Ahn, Rebecca L. Bond, Peter Dayan, Jonathan P. Roiser, and Oliver J. Robinson. 2019. "Altered Learning under Uncertainty in Unmedicated Mood and Anxiety Disorders." *Nature Human Behaviour* 3 (10): 1116–23.
- Badre, David, and Michael J. Frank. 2012. "Mechanisms of Hierarchical Reinforcement Learning in Cortico-Striatal Circuits 2: Evidence from fMRI." *Cerebral Cortex* 22 (3): 527–36.
- Baker, Daniel H., Greta Vilidaite, Freya A. Lygo, Anika K. Smith, Tessa R. Flack, Andre D. Gouws, and Timothy J. Andrews. 2019. "Power Contours: Optimising Sample Size and Precision in Experimental Psychology and Human Neuroscience." *arXiv [q-bio.NC]*. arXiv. <http://arxiv.org/abs/1902.06122>.
- Boehm, Udo, Helen Steingroever, and Eric-Jan Wagenmakers. 2018. "Using Bayesian Regression to Test Hypotheses about Relationships between Parameters and Covariates in Cognitive Models." *Behavior Research Methods* 50 (3): 1248–69.
- Brown, Vanessa M., Jiazhou Chen, Claire M. Gillan, and Rebecca B. Price. 2020. "Improving the Reliability of Computational Analyses: Model-Based Planning and Its Relationship with Compulsivity." *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, January. <https://doi.org/10.1016/j.bpsc.2019.12.019>.
- Clark, David M., Lauren Canvin, John Green, Richard Layard, Stephen Pilling, and Magdalena Janecka. 2018. "Transparency about the Outcomes of Mental Health Services (IAPT Approach): An Analysis of Public Data." *The Lancet* 391 (10121): 679–86.
- Daw, Nathaniel D. 2011. "Trial-by-Trial Data Analysis Using Computational Models." *Decision Making, Affect, and Learning: Attention and Performance XXIII* 23 (1). https://books.google.com/books?hl=en&lr=&id=IAHKGtbeunwC&oi=fnd&pg=PA3&ots=Jewelv8jp1&sig=3VKXGuZ8uSvFDXSHjtzwh_7TAMI.
- Dayan, Peter, and Nathaniel D. Daw. 2008. "Decision Theory, Reinforcement Learning, and the Brain." *Cognitive, Affective & Behavioral Neuroscience* 8 (4): 429–53.
- Gabry, Jonah, Daniel Simpson, Aki Vehtari, Michael Betancourt, and Andrew Gelman. 2019. "Visualization in Bayesian Workflow." *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 182 (2): 389–402.
- Gelman, Andrew, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. 2013. *Bayesian Data Analysis, Third Edition*. CRC Press.
- Huys, Quentin J. M., Tiago V. Maia, and Michael J. Frank. 2016. "Computational Psychiatry as a Bridge from Neuroscience to Clinical Applications." *Nature Neuroscience* 19 (3): 404–13.
- Kruschke, John. 2014. *Doing Bayesian Data Analysis: A Tutorial with R, JAGS, and Stan*. Academic Press.
- Matlab, Starting. 2012. "Matlab." *The MathWorks, Natick, MA*. https://itb.biologie.hu-berlin.de/~kempter/Teaching/2003_SS/gettingstarted.pdf.
- Mkrtchian, Anahit, Jessica Aylward, Peter Dayan, Jonathan P. Roiser, and Oliver J. Robinson. 2017. "Modeling Avoidance in Mood and Anxiety Disorders Using Reinforcement Learning." *Biological*

- Psychiatry* 82 (7): 532–39.
- Moutoussis, Michael, Alexandra K. Hopkins, and Raymond J. Dolan. 2018. "Hypotheses About the Relationship of Cognition With Psychopathology Should Be Tested by Embedding Them Into Empirical Priors." *Frontiers in Psychology* 9 (December): 2504.
- Nord, C. L., V. Valton, J. Wood, and J. P. Roiser. 2017. "Power-up: A Reanalysis of 'power Failure.'" *Neuroscience Using Mixture Modelling*, *Journal of Neuroscience*, 3592–3516.
- Penny, Will, Stefan Kiebel, and Karl Friston. 2003. "Variational Bayesian Inference for fMRI Time Series." *NeuroImage* 19 (3): 727–41.
- Ratcliff, Roger, Philip L. Smith, Scott D. Brown, and Gail McKoon. 2016. "Diffusion Decision Model: Current Issues and History." *Trends in Cognitive Sciences* 20 (4): 260–81.
- Robinson, Oliver J., and Henry W. Chase. 2017. "Learning and Choice in Mood Disorders: Searching for the Computational Parameters of Anhedonia." *Computational Psychiatry (Cambridge, Mass.)* 1 (1): 208–33.
- Scannell, Jack W., and Jim Bosley. 2016. "When Quality Beats Quantity: Decision Theory, Drug Discovery, and the Reproducibility Crisis." *PloS One* 11 (2): e0147215.
- Schonberg, Tom, Craig R. Fox, and Russell A. Poldrack. 2011. "Mind the Gap: Bridging Economic and Naturalistic Risk-Taking with Cognitive Neuroscience." *Trends in Cognitive Sciences* 15 (1): 11–19.
- Seymour, Ben, Nathaniel D. Daw, Jonathan P. Roiser, Peter Dayan, and Ray Dolan. 2012. "Serotonin Selectively Modulates Reward Value in Human Decision-Making." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 32 (17): 5833–42.
- Wilson, Robert C., and Anne Ge Collins. 2019. "Ten Simple Rules for the Computational Modeling of Behavioral Data." *eLife* 8 (November). <https://doi.org/10.7554/eLife.49547>.
- Wise, Toby, and Raymond J. Dolan. 2020. "Associations between Aversive Learning Processes and Transdiagnostic Psychiatric Symptoms Revealed by Large-Scale Phenotyping." *bioRxiv*. <https://doi.org/10.1101/843045>.
- Wise, T., J. Radua, E. Via, N. Cardoner, O. Abe, T. M. Adams, F. Amico, et al. 2017. "Common and Distinct Patterns of Grey-Matter Volume Alteration in Major Depression and Bipolar Disorder: Evidence from Voxel-Based Meta-Analysis." *Molecular Psychiatry* 22 (10): 1455–63.