

# MPTinR: Analysis of Multinomial Processing Tree Models in R\*

Henrik Singmann  
University of Zurich

Albert-Ludwigs-Universität Freiburg

David Kellen  
University of Basel

Albert-Ludwigs-Universität Freiburg

June 23, 2018

## Abstract

We introduce MPTinR, a software package developed for the analysis of multinomial processing tree (MPT) models. MPT models represent a prominent class of cognitive measurement models for categorical data with applications in a wide variety of fields. MPTinR is the first software for the analysis of MPT models in the statistical programming language R, providing a modeling framework that is more flexible than standalone software packages. MPTinR also introduces important features such as a) the ability to calculate the Fisher information approximation measure of model complexity for MPT models, b) the ability to fit models for categorical data outside the MPT model class, like signal-detection models, c) a function for model selection across a set of nested and non-nested candidate models (using several model selection indices), and d) multi-core fitting. MPTinR is available from the Comprehensive R Archive Network at <http://cran.r-project.org/web/packages/MPTinR/>

Keywords: MPT models, R, software, Fisher information

## 1 Introduction

Multinomial processing tree (MPT) models represent a prominent class of cognitive measurement models for categorical data (Hu and Batchelder, 1994; Purdy and Batchelder, 2009; Riefer and Batchelder, 1988). They describe the observed response frequencies from a finite set of response categories (i.e., responses following a multinomial distribution) with a finite number of latent states. Each latent state is reached by particular combinations of cognitive processes; processes that are assumed to take place in an all-or-nothing fashion (e.g., either a previously seen item is remembered as having been seen or not). The probability of a latent state being reached depends on the probabilities that the different cognitive processes associated to it successfully take place. The latent states usually follow each other in a serial order that can be displayed in a tree-like structure (see Figure 1, this model is described in more detail below).

MPT models exist in a wide range of fields, such as memory, reasoning, perception, categorization, and attitude measurement (for reviews see Erdfelder *et al.*, 2009; Batchelder and Riefer, 1999) where they provide a superior data analysis strategy compared to the usually employed ad-hoc models (e.g., ANOVA). MPT models allow the decomposition of observed responses into latent cognitive processes with a psychological interpretation, whereas ad-hoc models only permit the testing of hypotheses concerning the observed data, and the model parameters do not have a psychological interpretation nor does the model provide any insight into the underlying cognitive processes. In this paper we present a package for the analysis of MPT models for the statistical programming language R (R Development Core Team, 2012a) called **MPTinR** that offers several advantages over comparable software (see Moshagen, 2010, for a comparison of software for the analysis of MPTs).

The remainder of this paper is organized as follows. In the next section we will introduce a particular MPT model as an example. In the section thereafter we will provide a general overview of model representation,

---

\*This is an extended and updated version of our manuscript on MPTinR. The code given in this document can be executed as is and the output given here is produced from a real R session (using **knitr** and **sweave**). The original paper was published as: Singmann, H., & Kellen, D. (2013). MPTinR: Analysis of multinomial processing tree models in R. *Behavior Research Methods*, 45(2), 560-575. <http://doi.org/10.3758/s13428-012-0259-0>

parameter estimation, and statistical inference in the MPT model class. This overview is by no means exhaustive, but it gives several references that provide in-depth analysis. In the section thereafter, we introduce MPTinR and its functionalities and provide an example session introducing the most important functions for model fitting, model selection and simulation. Furthermore, an overview of all functions in MPTinR is given. Finally, the Appendix contains a description of the algorithms used by MPTinR.

## 2 An example MPT: The Two-High-Threshold Model of Recognition Memory

Consider the model depicted in Figure 1, which describes the responses produced in a simple recognition memory experiment consisting of two phases: A learning-phase in which participants study a list of items (e.g., words), and subsequently a test-phase in which a second list is presented and participants have to indicate which items were previously studied (old items) and which were not (new items) by responding “Old” or “New”, respectively.

This particular MPT model for the recognition task — the Two-High Threshold Model (2HTM; Snodgrass and Corwin, 1988) — has been chosen because of its simplicity. It consists of two trees, with the item-type associated with each tree (old and new items) specified at the tree’s root. Response categories are specified at the leaves of the trees. Cognitive processes are specified in a sequential manner by the tree nodes and their outcomes (successful occurrence or not) are represented by the branches that emerge from these nodes. The probability of each cognitive process successfully occurring is defined by a parameter.

Let us first consider the old-item tree. When presented with an old item at test, a state of successful remembering is reached with probability  $D_o$  (= detect old), and the “Old” response is then invariably given. If the item is not remembered (with probability  $(1 - D_o)$ ), the item is guessed to be “Old” with probability  $g$  (= guessing), or to be “New” with probability  $(1 - g)$ . Regarding the new-item tree, items can be detected as non-studied with probability  $D_n$  (= detect new), leading to the items’ rejection (response “New”). If the new item is not detected (with probability  $(1 - D_n)$ ), then it is guessed to be “Old” or “New” with probabilities  $g$  and  $(1 - g)$ , respectively. The predicted response probabilities for each observable response category can be represented by a set of equations:

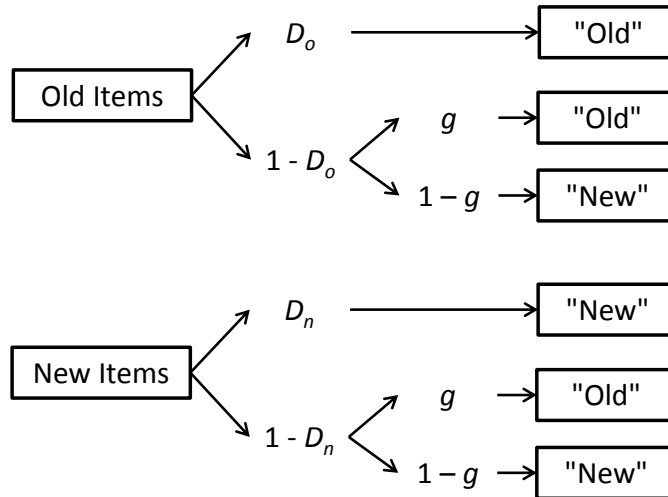


Figure 1: A MPT model (2HTM) for recognition memory. On the left side are the two different item types, old and new items, respectively, each represented by one tree. On the right side are the observed responses “Old” and “New”. In between are the assumed latent states with the probabilities leading to these states. Each tree is traversed from left to right.  $D_o$  = Detect an old item as old,  $D_n$  = detect a new item as new,  $g$  = guess an item as old.

$$P(\text{"Old"}|\text{old item}) = D_o + (1 - D_o) \times g \quad (1)$$

$$P(\text{"New"}|\text{old item}) = (1 - D_o) \times (1 - g) \quad (2)$$

$$P(\text{"Old"}|\text{new item}) = (1 - D_n) \times g \quad (3)$$

$$P(\text{"New"}|\text{new item}) = D_n + (1 - D_n) \times (1 - g) \quad (4)$$

These equations are constructed by concatenating all branches leading to the same observable response category (e.g., “Old”) within one tree. For example, the first line concatenates all branches leading to response “Old” in the old-item tree. As stated above, for old items the response “Old” is given either when an item is successfully remembered as being old ( $D_o$ ) or, if it is not remembered as being old, by guessing ( $(1 - D_o) \times g$ ). Note that the responses associated to the equations above only provide two degrees of freedom, while the model equations assume three free parameters ( $D_o$ ,  $D_n$ , and  $g$ ), which means that the model is in the present form non-identifiable (see Bamber and van Santen, 1985). This issue will be discussed in greater detail below.

The model presented above describes observed responses in terms of a set of unobservable latent cognitive processes, namely a mixture of 1) memory retrieval ( $D_o$ ), 2) distractor detection ( $D_n$ ), and 3) guessing ( $g$ ). Whereas the memory parameters are specific for the item types (i.e.,  $D_o$  is only part of the old item tree and  $D_n$  is only part of the new item tree), the same guessing parameter is present in both trees. This means that it is assumed that guessing (or response bias) is identical whether or not an item is old or new reflecting that the status of each item (old or new) is completely unknown to the participants when guessing. Note that this psychological interpretation of the parameters requires validation studies. In these studies it needs to be shown that certain experimental manipulations expected to selectively affect certain psychological processes are reflected in the resulting model parameters, with changes only being reliably found in the parameters representing those same processes (see Snodgrass and Corwin, 1988, for validation studies of 2HTM parameters).

The contribution of each of the assumed cognitive processes can be assessed by finding the parameter values (numerically or analytically) that produce the minimal discrepancies between predicted and observed responses. The discrepancies between predicted and observed responses can be quantified by a divergence statistic (Read and Cressie, 1988). As we discuss in more detail below, discrepancies between models and data can be used to evaluate the overall adequacy of the model and to test focused hypotheses on parameters (e.g., parameters have the same values across conditions).

### 3 Representation, Estimation and Inference in MPT Models: A Brief Overview

#### 3.1 Model Specification and Parameter Estimation

Following the usual formalization (Hu and Batchelder, 1994) of MPT models, let  $\Theta = \{\theta_1, \dots, \theta_S\}$ , with  $0 \leq \theta_s \leq 1$ ,  $\theta_s$ ,  $s = 1, \dots, S$  denote the vector of  $S$  model parameters representing the different cognitive processes. For tree  $k$ , the probability of a branch  $i$  leading to response category  $j$  given  $\Theta$  corresponds to:

$$p_{i,j,k}(\Theta) = c_{i,j,k} \prod_{s=1}^S \theta_s^{a_{s,i,j,k}} (1 - \theta_s)^{b_{s,i,j,k}}, \quad (5)$$

where  $a_{s,i,j,k}$  and  $b_{s,i,j,k}$  represent the number of times each parameter  $\theta_s$  and its complement ( $1 - \theta_s$ ) are respectively represented at each branch  $i$  leading to category  $j$  of tree  $k$ , and  $c_{i,j,k}$  represents the product of constants on the tree links, if the latter are present. The probability of category  $j, k$  given  $\Theta$  corresponds to  $p_{j,k}(\Theta) = \sum_{i=1}^{J_k} p_{i,j,k}(\Theta)$  (i.e., the sum of all branches ending in one response category per tree), with  $\sum_{j=1}^{J_k} p_{j,k}(\Theta) = 1$  (i.e., the sum of all probabilities per tree is 1).

Let  $\mathbf{n}$  be a vector of observed category frequencies, with  $n_{j,k}$  denoting the frequency of response category  $j$  in tree  $k$ , with  $N_k = \sum_{j=1}^{J_k} n_{j,k}$  and  $N = \sum_{k=1}^K N_k$ . The likelihood function of  $\mathbf{n}$  given model parameter

vector  $\Theta$  is:

$$L(\mathbf{n} | \Theta) = \prod_{k=1}^K \left[ \binom{N_k}{n_{1,k}, \dots, n_{J_k,k}} \prod_{j=1}^{J_k} [p_{j,k}(\Theta)^{n_{j,k}}] \right] \quad (6)$$

The parameter values that best describe the observed responses correspond to the ones that maximize the likelihood function in Equation 6. These maximum-likelihood parameter estimates (denoted by  $\hat{\Theta}$ ) can sometimes be obtained analytically (e.g., Stahl and Klauer, 2008), but in the vast majority of cases they can only be found by means of iterative methods such as the EM algorithm (see Hu and Batchelder, 1994). Regarding the variability of the maximum-likelihood parameter estimates, confidence intervals can be obtained by means of the Fisher Information Matrix (the matrix of second order partial derivatives of the likelihood function with respect to  $\Theta$ ; Riefer and Batchelder, 1988) or via bootstrap simulation (Efron and Tibshirani, 1994).

The search for the parameters that best describe the data (maximize the likelihood function) requires that the model is *identifiable*. Model identifiability concerns the property that a set of predicted response probabilities can only be obtained by a single set of parameter values. Let  $\Theta$  and  $\Theta'$  be model parameter vectors, with  $p(\Theta)$  and  $p(\Theta')$  as their respective predicted response probabilities. A model is globally identifiable if  $\Theta \neq \Theta'$  implies  $p(\Theta) \neq p(\Theta')$  across the entire parameter space, and locally identifiable if it holds in the region of parameter space where  $\hat{\Theta}$  lies (Schmittmann *et al.*, 2010). An important aspect is that the degrees-of-freedom provided by a dataset provide the upper bound for the number of potentially identifiable free parameters in an model, that is  $S \leq \sum_{k=1}^K (J_k - 1)$ . Local identifiability is sufficient for most purposes, and it can be shown to hold by checking that the Fisher Information Matrix for  $\hat{\Theta}$  has rank equal to the number of free parameters (the rank of the Fisher Information Matrix is part of the standard output produced by MPTinR). For a detailed discussion on model identifiability in the MPT model class, see Schmittmann *et al.* (2010).

### 3.2 Null-Hypothesis Testing

The discrepancies between predicted and observed response frequencies when taking the maximum-likelihood parameter estimates ( $\hat{\Theta}$ ) are usually summarized by the  $G^2$  statistic:

$$G^2 = 2 \sum_{k=1}^K \sum_{j=1}^{J_k} n_{j,k} [\ln(n_{j,k}) - \ln(N_k p_{j,k})] \quad (7)$$

The smaller  $G^2$ , the smaller the discrepancies.<sup>1</sup> An important aspect of the  $G^2$  statistic is that it follows a chi-square distribution with degrees-of-freedom equal to the number of independent response categories minus the number of free parameters  $\left( \left[ \sum_{k=1}^K (J_k - 1) \right] - S \right)$ . This means that the quality of the account provided by the model can be assessed through null-hypothesis testing. Parameter-equality restrictions (e.g.,  $\theta_1 = \theta_2$  or  $\theta_1 = 0.5$ ) can also be tested by means of null-hypothesis testing. The difference in  $G^2$  between the unrestricted and restricted models also follows a chi-square distribution with degrees-of-freedom equal to the difference in free parameters between the two models (Riefer and Batchelder, 1988). It should be noted

---

<sup>1</sup>Parameter estimation in MPTinR can only be made using the maximum-likelihood method (Equation 6), which can be obtained by minimizing the discrepancies between observed and expected response frequencies as measured by the  $G^2$  statistic (Equation 7; Bishop *et al.*, 1975). Instead of minimizing the  $G^2$  statistic, other discrepancy statistics can be used, in particular one of the many possible statistics coming from the power-divergence family (Read and Cressie, 1988), from which the  $G^2$  is a special case. Studies (e.g., García-Pérez, 1994; Riefer and Batchelder, 1991) have shown that some of these statistics can be advantageous when dealing with sparse data and attempting to minimize a model's sensitivity to outliers. Still, the use of alternatives to the  $G^2$  statistic coming from the power-divergence family has several shortcomings: First, it is not clear which particular statistic would be more advantageous for a specific type of MPT model and data, a situation that would require an extensive evaluation of different alternative statistics. Second, the use of an alternative to the  $G^2$  statistic represents a dismissal of the maximum-likelihood method, which in turn compromises the use of popular model selection measures such as the Akaike Information Criterion, the Bayesian Information Criterion, or the Fisher Information Approximation (which will be discussed in detail later), all of which assume the use of the maximum-likelihood method. Given these disadvantages, and the almost-ubiquitous use of maximum-likelihood estimation in MPT modeling (for reviews, see Batchelder and Riefer, 1999; Erdfelder *et al.*, 2009) the current version of MPTinR only implements the maximum-likelihood method for parameter estimation and the  $G^2$  statistic for quantification of model misfit.

that parameter-inequality restrictions (e.g.,  $\theta_1 \leq \theta_2$ ; see Knapp & Batchelder, 2004) can also be tested, but the difference in  $G^2$  no longer follows a chi-square distribution but a particular mixture of chi-square distributions that in many cases needs to be determined via simulation (Silvapulle and Sen, 2005, for an example see Kellen and Klauer, 2011).

### 3.3 Model Selection

It is important to note that  $G^2$  is a measure that only summarizes the models' goodness-of-fit, and which can only be used to test between nested models (when the restricted model is a special case of the unrestricted model). Furthermore, it ignores possible flexibility differences between the models, that is differences in their inherent ability to fit data in general. The more flexible a model is, the better it will fit any data pattern, regardless of the appropriateness of the model. The best model is not necessarily the one that better fits the data, as it is also important that a model's range of predictions closely follows the observations made and that it can produce accurate predictions regarding future observations (Roberts and Pashler, 2000). Model selection analyses attempt to find the model that strikes the best balance between goodness-of-fit and model flexibility (for discussions on different model selection approaches, see Myung *et al.*, 2000; Wagenmakers and Waldorp, 2006), which makes  $G^2$  an unsuitable measure for the comparison of non-nested models. In order to compare both nested and non-nested models in a single framework, as well as to account for potential differences in model flexibility, measures such as the Akaike Information Criterion (AIC; Akaike, 1974) and the Bayesian Information Criterion (BIC; Schwarz, 1978) are used:<sup>2</sup>

$$\text{AIC} = G^2 + 2S \tag{8}$$

$$\text{BIC} = G^2 + \ln(N)S \tag{9}$$

AIC and BIC correct models' fit-results by introducing a punishment factor (the second term in the formulas) that penalizes them for their flexibility ( $S$  is the number of parameters). The lower the AIC/BIC the better the account. For the case of AIC and BIC, the number of free parameters is used as a proxy for model flexibility, a solution that is convenient to use but that ignores differences in the model's functional form and is rendered useless when used to compare models that have the same number of parameters (Klauer and Kellen, 2011). For example, consider the structurally identical models A and B with two parameters  $\theta_1$  and  $\theta_2$ , with the sole difference between both models that for model B the restriction  $\theta_1 \leq \theta_2$  holds. According to AIC and BIC the models are equally flexible despite the fact that the inequality restriction halves model B's parameter space and therefore its flexibility.

A measure that provides a more precise quantification of model flexibility is the Fisher Information Approximation (FIA), a measure that stems from the Minimum Description Length framework (for an introduction, see Grünwald, 2007):

$$\text{FIA} = \frac{1}{2}G^2 + \frac{S}{2}\ln\frac{N}{2\pi} + \ln \int \sqrt{\det I(\Theta)} d\Theta \tag{10}$$

where  $I(\Theta)$  is the Fisher Information Matrix for sample size 1 (for details, see Su *et al.*, 2005). The third additive term of Equation 10 is the penalty factor that accounts for the functional form of the model, providing a more accurate depiction of a model's flexibility. Unlike AIC and BIC, FIA can account for flexibility differences in models that have the same number of parameters. Despite its advantages, FIA is a measure whose computation is far from trivial given the integration of the determinant of the Fisher Information Matrix across a multidimensional parameter space. Due to the recent efforts of Wu *et al.* (2010a,b) the computation of FIA for the MPT model class has become more accessible.

---

<sup>2</sup>In the AIC and BIC formulas, the first term corresponds to the model's goodness of fit, and the second additive term to the model's penalty factor. As noted by one of the reviewers, we use the  $G^2$  as the first term, contrary to other implementations that use the models' log-likelihood ( $LL_M$ ) instead.  $G^2$  corresponds to  $2 \times (LL_S - LL_M)$ , with  $LL_S$  being the log-likelihood of a saturated model that perfectly describes the data. In this sense, the definitions of AIC and BIC given in the main body of text can be viewed as differences in AIC and BIC between these two models, making the notation  $\Delta\text{AIC}$  and  $\Delta\text{BIC}$  more appropriate. We nevertheless use the notation AIC and BIC given that we use the notation  $\Delta\text{AIC}$  and  $\Delta\text{BIC}$  when referring to differences between different candidate models other than the saturated model that perfectly describes the data.

### 3.4 Context-Free Language for MPTs

Also of interest is the context-free language for the MPT model class developed by Purdy and Batchelder (2009), called  $L_{\text{BMPT}}$ . In  $L_{\text{BMPT}}$ , each MPT model is represented by a string, called a word, consisting only of symbols representing parameters ( $\theta$ ) or categories ( $C$ ). The word in  $L_{\text{BMPT}}$  representing each tree is created by recursively performing the following operations:

1. visit the root
2. traverse the upper subtree
3. traverse the lower subtree

During these operations the word is built in the following manner: Whenever a parameter (and *not* its converse) is encountered, add the parameter to the string. Whenever a response category (i.e., leaf) is reached, add the category to the string. The word is complete when reaching the last response category. The structure for the trees in Figure 1 in  $L_{\text{BMPT}}$  is thus:

$$\theta C \theta C C \tag{11}$$

By assigning indices one obtains a word in  $L_{\text{BMPT}}$  for each tree in Figure 1:

$$\begin{aligned} &\theta_{D_o} C_{Old} \theta_g C_{Old} C_{New} \\ &\theta_{D_n} C_{New} \theta_g C_{Old} C_{New} \end{aligned} \tag{12}$$

In order to create a single MPT model of the two trees in Figure 1 one needs to assume a joining parameter  $\theta_{join}$  whose branches connect the two trees into a single one. In this case, the values of  $\theta_{join}$  and  $(1-\theta_{join})$  would be fixed a priori as they represent the proportion of times that old and new items occur during test. The resulting full model for the recognition memory experiment in  $L_{\text{BMPT}}$  is:

$$\theta_{join} \theta_{D_o} C_{Old} \theta_g C_{Old} C_{New} \theta_{D_n} C_{New} \theta_g C_{Old} C_{New} \tag{13}$$

The context-free language of Purdy and Batchelder (2009) is extremely useful as it allows the statement and proof of propositions regarding the MPT model class. One application of this language is in the computation of FIA (Wu et al., 2010a, 2010b).

## 4 General Overview of MPTinR

**MPTinR** offers five main advantages over comparable software (cf. Moshagen, 2010). First, **MPTinR** is fully integrated into the R language, an open-source implementation of the S statistical programming language (Becker *et al.*, 1988), that is becoming the lingua franca of statistics (Muenchen, 2012; Vance, 2009). Besides being free (as it is part of the GNU project, see <http://www.gnu.org>) and platform independent, R's major strength is the combination of being extremely powerful (it is a programming language) with the availability of a wide variety of statistical and graphical techniques. Couching **MPTinR** within this environment lets it benefit from these strengths. For example, data usually needs to be preprocessed before fitting an MPT model. In addition to fitting an MPT model, one may want to visualize the parameter estimates, run hypothesis-tests on particular parameter restrictions or perform simulations validating certain aspects of the model such as the parameter estimates or the identifiability of the model. When using **MPTinR** all of those processes that can be done within one single environment without the need to ever move data between programs.

Second, **MPTinR** was developed with the purpose of being easy to use, improving some of the more cumbersome features of previous programs, such as the ones concerning model representation. MPT models are represented in most programs such as GPT (Hu and Phillips, 1999), HMMTree (Stahl and Klauer, 2007), or multiTree (Moshagen, 2010) by means of .EQN model files. Model specification in .EQN files need to follow a certain structure that could lead to errors, and diverge from the model equations (e.g., Equations 1-4) that are normally used to represent these models in scientific manuscripts. These requirements can become especially cumbersome when handling MPT models comprised of trees with numerous branches

(e.g., Oberauer, 2006). Furthermore, most programs require parameter restrictions to be specified “by hand” every time the program is used (for an exception, see Moshagen, 2010), or different model files implementing the parameter restrictions have to be created. **MPTinR** overcomes these inconvenient features: Models are specified in a way that is virtually equivalent to the equations used to represent models (i.e., Equations 1-4), and model restrictions are intuitively specified. Furthermore, model and restrictions can be specified in external files or directly within an R script. In addition, **MPTinR** automatically detects whether single or multiple datasets are fitted and adjusts the output accordingly. For multiple datasets, the summed results (e.g. summed  $G^2$  values) as well as the results for the aggregate data (i.e., summed response frequencies across datasets) are per default provided in the output.

Third, **MPTinR** provides different model selection measures, namely AIC, BIC, and FIA. As previously referred, the computation of FIA is not trivial, and only very recently has it become available for the MPT model class (Wu *et al.*, 2010a,b).

Fourth, **MPTinR** is able to translate an MPT model into a string representation according to the context-free language developed by Purdy and Batchelder (2009) as well as construct model equations from the string representation. Given that the manual translation of MPTs can be rather difficult and tedious, the possibility of an automatic translation will likely encourage the use of this context-free language, which has shown great potential in the assessment of model flexibility (see Purdy, 2011).

Fifth, although being specifically designed for MPTs, **MPTinR** can also be used to fit a wide range of other cognitive models for categorical data, for example models based on signal detection theory (SDT; Green and Swets, 1966; Macmillan and Creelman, 2005). This essentially makes **MPTinR** a framework for fitting many types of cognitive models for categorical data and to facilitate their comparison. As this last point is outside the scope of this article we refer interested readers to the documentation for the functions `fit.model` and `fit.mptinr` which contain detailed examples of how to fit different SDT models.

## 4.1 Getting Started

**MPTinR** is a package for the R programming language and therefore needs to be used within the R environment by using the functions described below. For users familiar with commercial statistic packages such as SPSS, the handling of R may be uncommon as it does not come with a graphical user interface (but see Valero-Mora and Ledesma (2012), for an overview, and Rödiger *et al.* (2012), for a powerful graphical user interface for R). Instead, all commands have to be entered at the prompt. **MPTinR** comes with a manual describing all functions in detail (available also via <http://cran.r-project.org/web/packages/MPTinR/MPTinR.pdf>) and has a website with more information on important features such as model files and restrictions (see <http://www.psychologie.uni-freiburg.de/Members/singmann/R/mptinr/modelfile>). To obtain the documentation for any function of **MPTinR** simply enter the function name at the prompt preceded by a ? (e.g., `?fit.mpt` to obtain the detailed documentation containing examples for the main function `fit.mpt`). The documentation for each function contains a detailed description of its use and the arguments that need to be passed to it. Here we only present the most relevant arguments of each function.

**MPTinR** is available via the Comprehensive R Archive Network (CRAN; <http://cran.r-project.org/>) and can therefore be installed from within any R session with the following command (given an active Internet connection):

```
install.packages("MPTinR")
```

Note that you might need an up-to-date version of R to install **MPTinR**. After successful installation (which only needs to be done once), **MPTinR** needs to be loaded into the current R session with the following command (this needs to be done each time you start a new R session):

```
library("MPTinR")
```

## 4.2 Format of Models, Restrictions, and Data

The basis of all analyses of MPT models in **MPTinR** is the representation of the model via a model file. Whereas **MPTinR** can read the well-known .EQN model files (e.g., Hu and Phillips, 1999) it offers an alternative, the **easy** format. To specify a model in the **easy** format the model file needs to contain the right

hand sides of the equations defining an MPT model (e.g., Equations 1-4) with the equations for each tree separated by at least a single empty line. In other words, for each tree all branches ending in the same response category need to be written in a single line concatenated by +. Note that only trees with binary branching can be specified in **MPTinR** (for an exception, see Hu and Phillips (1999)). The model file for the 2HTM model from Figure 1 in the **easy** format could be:

```
# Tree for old items: First 'yes', then 'no'
Do + (1 - Do) * g
(1-Do)*(1-g)

#Tree for new items: First 'yes', then 'no'
(1-Dn) * g
Dn + (1-Dn) * (1 - g)
```

As can be seen, MPTinR allows for comments in the model file. Everything to the right of the number sign # will be ignored and lines containing only a comment count as empty.<sup>3</sup> Also, additional whitespace within the equations is ignored. Note that the parameter names used in the model files need to be valid R variable names (for details, type `?make.names` and `?reserved` at the command prompt).

The format of restrictions files is similar to the format of model files. Each restriction needs to be specified on one line. Furthermore, the following rules apply: Inequality restrictions needs to be placed before equality restrictions and can only be specified using the smaller than operator < (note that inequality restrictions containing < actually represent the weak inequality  $\leq$ ). If a variable appears not as the rightmost element in a restriction it can only appear as the rightmost element in any other restriction (in other words in a set of restrictions a variable can appear multiple times, but only once not as the rightmost element). In addition to simple equality and inequality restrictions MPTinR can also deal with order restrictions involving more than two parameters. For example,  $Y1 = Y2 = 0.5$  will set both parameters  $Y1$  and  $Y2$  to 0.5. Similarly,  $W1 < W2 < W3$  will be correctly interpreted as  $W1 < W2$  and  $W2 < W3$ . A valid restrictions file (for a fictitious MPT model) could be (note that we have added quotes at the beginning and end for display purposes, those need to be removed in an actual file):

```
" # quotes need to be removed
W1 < W2 < W3
X4 = X3
Y1 = Y3 = 0.5
Z = 0 #Restrictions may also contain comments.
"
```

Note that it is also possible to specify model and restriction within an R script (as compared to in an external file), using for example the `textConnection` function included in the base R package. Restrictions can also be specified within a script as a list of strings. We will use this functionality in the following.

**MPTinR** contains the function `check.mpt` that may help in writing model and restrictions files. It has the format `check.mpt(model.filename, restrictions.filename = NULL)` and will return a list with the following information: a logical value indicating whether or not the probabilities on each tree sum to one, the number of trees and the number of response categories the model has, the number of independent response categories the model provides, and the name and number of the parameters in the tree. `model.filename` is the only mandatory argument and needs to be a character string specifying the location and name of the model file<sup>4</sup>. `restrictions.filename` is an optional argument specifying the location and name of the restrictions file.

<sup>3</sup>Note that the way MPTinR deals with comments has changed. As is common in R (and other programming languages) everything to the right of the comment symbol # is ignored. In previous versions of MPTinR (prior to version 0.9.2) a line containing a # at any position was ignored completely.

<sup>4</sup>R looks for files in the current working directory. To find out what is the current working directory type `getwd()` at the R prompt. You can change the working directory using either the R menu or using function `setwd`. Additionally, models and restrictions can also be specified within an R script (i.e., not in a file) using a `textConnection`, see the examples in `?fit.mpt` and `?fit.model`.



For example, calling `check.mpt` on the 2HTM model above will return an output indicating that the probabilities in each tree sum to one (if not, the function will pinpoint the misspecified trees), that the number of trees in the model is two, the number of response categories is four, the number of independent response categories is two, and the three parameters are  $D_n$ ,  $D_o$ , and  $g$ :

```
mod_2htm_1 <- "
# Tree for old items: First 'yes', then 'no'
Do + (1 - Do) * g
(1-Do)*(1-g)

#Tree for new items: First 'yes', then 'no'
(1-Dn) * g
Dn + (1-Dn) * (1 - g)
"

check.mpt(textConnection(mod_2htm_1))

## $probabilities.eq.1
## [1] TRUE
##
## $n.trees
## [1] 2
##
## $n.model.categories
## [1] 4
##
## $n.independent.categories
## [1] 2
##
## $n.params
## [1] 3
##
## $parameters
## [1] "Dn" "Do" "g"
```

The data on which an MPT model will be fitted needs to be passed as a numeric data object, either as a `vector`, `matrix`, or `data.frame`. The mapping of data to response category is done via position. The ordinal position (rank) of each equation in the model file needs to correspond to the response category that is represented at that position/rank in the data `vector` or to the column with that number if the data object is a `matrix` or `data.frame`. If a `matrix` or `data.frame` contains more than one row, each row is considered as one dataset and the MPT model is fitted separately for each dataset and the data summed across rows is considered as another dataset (called `aggregated` dataset) which is also fitted. The data can be entered directly into R or by loading it using one of the data import functions (e.g., `read.table`, see also R Development Core Team, 2012b). The `aggregated` data of the dataset described below could be entered as vector `d.broeder.agg` as:

```
d.broeder.agg <- c(145, 95, 170, 1990, 402, 198, 211, 1589, 868, 332,
                 275, 925, 1490, 310, 194, 406, 1861, 299, 94, 146)
```

### 4.3 An Example Session

Before estimating model parameters, it is important to see whether a model is identifiable. As previously pointed out the 2HTM as presented in Figure 1 has three parameters, while the “Old”/”New” responses for both old and new items only provide two independent categories (i.e., independent data points to be fitted) as given in the output of `check.mpt`. This means that the 2HTM with three parameters is not identifiable in the

current form. There are two ways of achieving identifiability for this model: 1) by imposing the restriction  $D_n = D_o$  (Snodgrass and Corwin, 1988), and/or 2) by including additional sets of observed categorical responses and subsequently extend the model in order to account for them.

The extension proposed by the second option can be implemented by fitting the model to responses obtained across different bias conditions, where individuals assumed distinct tendencies to respond “Old” or “New”. These different response biases or tendencies can be induced by changing the proportion of old items in the test phase (e.g., 10% vs. 90%), for example. According to the theoretical principles underlying the 2HTM, the guessing parameter  $g$  would be selectively affected by a response-bias manipulation, with  $D_o$  and  $D_n$  remaining unchanged. For example, Bröder and Schütz (2009) used the second solution sketched above by implementing five separate test phases, each with different proportions of old items (10%, 25%, 50%, 75%, and 90%). Consider the resulting model `mod_2htm_2` and the corresponding output from `check.mpt` showing that there are more degrees-of-freedom than free parameters:

```
mod_2htm_2 <- "
# Tree for old items (10%): First 'yes', then 'no'
Do + (1 - Do) * g1
(1-Do)*(1-g1)

#Tree for new items (90%): First 'yes', then 'no'
(1-Dn) * g1
Dn + (1-Dn) * (1 - g1)

# Tree for old items (25%): First 'yes', then 'no'
Do + (1 - Do) * g2
(1-Do)*(1-g2)

#Tree for new items (75%): First 'yes', then 'no'
(1-Dn) * g2
Dn + (1-Dn) * (1 - g2)

# Tree for old items (50%): First 'yes', then 'no'
Do + (1 - Do) * g3
(1-Do)*(1-g3)

#Tree for new items (50%): First 'yes', then 'no'
(1-Dn) * g3
Dn + (1-Dn) * (1 - g3)

# Tree for old items (75%): First 'yes', then 'no'
Do + (1 - Do) * g4
(1-Do)*(1-g4)

#Tree for new items (25%): First 'yes', then 'no'
(1-Dn) * g4
Dn + (1-Dn) * (1 - g4)

# Tree for old items (90%): First 'yes', then 'no'
Do + (1 - Do) * g5
(1-Do)*(1-g5)

#Tree for new items (10%): First 'yes', then 'no'
(1-Dn) * g5
Dn + (1-Dn) * (1 - g5)
"
```

```

check.mpt(textConnection(mod_2htm_2))

## $probabilities.eq.1
## [1] TRUE
##
## $n.trees
## [1] 10
##
## $n.model.categories
## [1] 20
##
## $n.independent.categories
## [1] 10
##
## $n.params
## [1] 7
##
## $parameters
## [1] "Dn" "Do" "g1" "g2" "g3" "g4" "g5"

```

Now, consider a  $40 \times 20$  matrix named `d.broeder` containing the individual data of the 40 participants from Bröder and Schütz (2009) Experiment 3.<sup>5</sup> Each participant was tested across five different base-rate conditions (10%, 25%, 50%, 75%, and 90% old items). In this data matrix, each row corresponds to one participant, and the columns correspond to the different response categories in the same order as in the model file. We just show the first 6 lines (i.e., participants) here:

```

data("d.broeder")
head(d.broeder)

##      [,1] [,2] [,3] [,4] [,5] [,6] [,7] [,8] [,9] [,10] [,11] [,12] [,13]
## [1,]   1   5   1  53  10   5   2  43  22   8  10  20  39
## [2,]   6   0   6  48  11   4   1  44  25   5   2  28  39
## [3,]   3   3   3  51  10   5   1  44  15  15   7  23  40
## [4,]   3   3   2  52  10   5   3  42  22   8   3  27  32
## [5,]   5   1   2  52  12   3  12  33  23   7  10  20  40
## [6,]   6   0   0  54  13   2   0  45  28   2   0  30  44
##      [,14] [,15] [,16] [,17] [,18] [,19] [,20]
## [1,]     6     5    10    49     5     4     2
## [2,]     6     2    13    48     6     1     5
## [3,]     5     3    12    42    12     2     4
## [4,]    13     1    14    39    15     3     3
## [5,]     5     9     6    51     3     6     0
## [6,]     1     0    15    49     5     0     6

```

**MPTinR** provides two main functions for model fitting and selection, `fit.mpt` and `select.mpt`:

`fit.mpt` is the major function for fitting MPT models to data, returning results such as the obtained log-likelihood and  $G^2$  value, the information criteria AIC, BIC, and FIA (if requested), parameter estimates and respective confidence intervals, and predicted response frequencies. Optionally one can specify model restrictions or request the estimation of the FIA. It has the basic format `fit.mpt(data, model.filename, restrictions.filename = NULL, n.optim = 5, fia = NULL)`. Two arguments in the `fit.mpt` function are of note: First, `fit.mpt` by default returns the best of five fitting runs for each dataset, a number that can be changed with the `n.optim` argument. Second, FIA is calculated using Monte Carlo methods (see Wu

<sup>5</sup>We thank Arndt Bröder for providing this dataset which also comes with **MPTinR**. The other files necessary to fit this data (i.e., the model and restriction files) also come with MPTinR, see `?fit.mpt`.

et al., 2010a, 2010b), with the number of samples to be used being specified by the `fia` argument. Following the recommendations of Wu et al., the number of samples should not be below 200,000 for real applications (we use smaller numbers in this vignette for demonstration purposes only).

`select.mpt` is a function that should aid in model selection (e.g., Zucchini, 2000) and takes multiple results from `fit.mpt` as the argument and produces a table comparing the models based on the information criteria AIC, BIC, and FIA. It has the basic format `select.mpt(mpt.results, output = c("standard", "full"))`, where `mpt.results` is a list of results returned by `fit.mpt`.

In order to exemplify the use of the `fit.mpt` and `select.mpt` functions, two MPT models are fitted to the data from Bröder and Schütz (2009) Experiment 3, the 2HTM model described above, and a restricted version of the 2HTM in which the  $g_1 < g_2 < g_3 < g_4 < g_5$  constraint was imposed<sup>6</sup>. As noted before, in real applications one should use more than 25,000 FIA samples (usually values such as 200,000).

```
br.2htm <- fit.mpt(d.broeder, textConnection(mod_2htm_2), fia = 25000)

## [1] "Computing FIA: Iteration begins at 2018-06-23 18:30:05"
## [1] "Computing FIA: Iteration stopped at 2018-06-23 18:30:06"
## Time difference of 0.7941077 secs

## Presenting the best result out of 5 minimization runs.

## [1] "Model fitting begins at 2018-06-23 18:30:06"

## Optimization routine for dataset(s) 2 6 7
## did not converge succesfully. Tried again with use.gradient == FALSE.
## Optimization for dataset(s) 2 6 7
## using numerically estimated gradients produced better results. Using those results.
## Old results saved in output == 'full' [['optim.runs']].

## [1] "Model fitting stopped at 2018-06-23 18:30:16"
## Time difference of 10.00993 secs
```

```
br.2htm.ineq <- fit.mpt(d.broeder, textConnection(mod_2htm_2),
                      list("g1 < g2 < g3 < g4 < g5"), fia = 25000)

## [1] "Computing FIA: Iteration begins at 2018-06-23 18:30:17"
## [1] "Computing FIA: Iteration stopped at 2018-06-23 18:31:01"
## Time difference of 44.61013 secs

## Presenting the best result out of 5 minimization runs.

## [1] "Model fitting begins at 2018-06-23 18:31:01"

## Optimization routine for dataset(s) 6 7
## did not converge succesfully. Tried again with use.gradient == FALSE.
## Optimization for dataset(s) 6 7
## using numerically estimated gradients produced better results. Using those results.
## Old results saved in output == 'full' [['optim.runs']].

## [1] "Model fitting stopped at 2018-06-23 18:31:21"
## Time difference of 19.37525 secs
```

As can be seen from the output, fitting a model produces various status messages and occasional warnings. The warnings stem from the numerical optimization algorithm employed (a quasi-Newton method, see Appendix for more details). Per default **MPTinR** uses the analytic gradient in the optimization process. If

<sup>6</sup>To fit only the aggregated data entered before as `d.broeder.agg` simply replace `d.broeder` with `d.broeder.agg` as in `br.2htm.2 <- fit.mpt(d.broeder.agg, "2htm.model")`.

the optimization algorithm indicates it did not converge successful, fitting is restarted but this time using the numerical gradient. The warning details which method produced a better result (i.e., larger maximum likelihood) and is consequently used. This warning usually indicates that data in some cells are missing (as is the case here) and some parameters might not be identified for those participants.

```
d.broeder[c(2, 6, 7),]

##      [,1] [,2] [,3] [,4] [,5] [,6] [,7] [,8] [,9] [,10] [,11] [,12] [,13]
## [1,]    6    0    6   48   11    4    1   44   25    5    2   28   39
## [2,]    6    0    0   54   13    2    0   45   28    2    0   30   44
## [3,]    6    0    0   54   13    2    3   42   28    2    1   29   44
##      [,14] [,15] [,16] [,17] [,18] [,19] [,20]
## [1,]     6     2    13   48     6     1     5
## [2,]     1     0    15   49     5     0     6
## [3,]     1     1    14   51     3     1     5
```

`fit.mpt` returns a list with the following elements (for a detailed description see `?fit.mpt`):

```
str(br.2htm, 1)

## List of 6
## $ goodness.of.fit      :List of 3
## $ information.criteria:List of 3
## $ model.info           :List of 2
## $ parameters           :List of 3
## $ data                 :List of 2
## $ fitting.runs        :List of 2
```

Element `model.info` can be used for further model diagnostic. As described before, a necessary condition that all parameters are identified is that the rank of the Fisher Matrix is equal to the number of free parameters. As expected from the warning during fitting, this is not the case for participants 2 and 7, while participant 6 showed no problems here.

```
br.2htm$model.info

## $individual
##      rank.fisher n.parameters n.independent.categories
## 1             7             7             10
## 2             6             7             10
## 3             7             7             10
## 4             7             7             10
## 5             7             7             10
## 6             7             7             10
## 7             6             7             10
## 8             7             7             10
## 9             7             7             10
## 10            7             7             10
## 11            7             7             10
## 12            7             7             10
## 13            7             7             10
## 14            7             7             10
## 15            7             7             10
## 16            7             7             10
## 17            7             7             10
## 18            7             7             10
```

```

## 19      7      7      10
## 20      7      7      10
## 21      7      7      10
## 22      7      7      10
## 23      7      7      10
## 24      7      7      10
## 25      7      7      10
## 26      7      7      10
## 27      7      7      10
## 28      7      7      10
## 29      7      7      10
## 30      7      7      10
## 31      7      7      10
## 32      7      7      10
## 33      7      7      10
## 34      7      7      10
## 35      7      7      10
## 36      7      7      10
## 37      7      7      10
## 38      7      7      10
## 39      7      7      10
## 40      7      7      10
##
## $aggregated
##   rank.fisher n.parameters n.independent.categories
## 1           7           7           10

```

Let us now look at the 2HTM results for aggregated data (which replicate the results from Bröder and Schütz, 2009, Table 4):

```

br.2htm[["goodness.of.fit"]][["aggregated"]]

##   Log.Likelihood G.Squared df   p.value
## 1      -5376.127  2.835736  3 0.4176509

br.2htm[["parameters"]][["aggregated"]]

##   estimates lower.conf upper.conf
## Dn 0.4450308 0.3716383 0.5184232
## Do 0.5561326 0.5102111 0.6020540
## g1 0.1411012 0.1129961 0.1692063
## g2 0.2152983 0.1736149 0.2569817
## g3 0.3998479 0.3388262 0.4608696
## g4 0.6045703 0.5446874 0.6644531
## g5 0.6895514 0.6424590 0.7366437

```

Next we may want to compare the parameter estimates for the aggregated data between the original and inequality restricted model:

```

br.2htm[["parameters"]][["aggregated"]][,"estimates"]

## [1] 0.4450308 0.5561326 0.1411012 0.2152983 0.3998479 0.6045703 0.6895514

br.2htm.ineq[["parameters"]][["aggregated"]][,"estimates"]

## [1] 0.4450308 0.5561326 0.1411012 0.2152983 0.3998479 0.6045703 0.6895514

```

As these results show, the order restriction on the guessing parameters held for the aggregated datasets (as the parameter values are identical between the two models) and the parameters values are within reasonable ranges. Note that order of the parameters is alphabetical (i.e.,  $D_n, D_o, g_1, g_2, g_3, g_4, g_5$ ). This ordering is based on the current locale (a fact that can lead to unexpected behaviors when e.g., moving between operating systems employing different locales such as Linux and Windows)! Before comparing the models to decide which to select based on the performance on this dataset we might want to check if all models provided a reasonable account of the data by inspecting the goodness of fit statistics. To this end we not only inspect the aggregated data but also the summed individual fits.

```
br.2htm[["goodness.of.fit"]][2:3]

## $sum
##   Log.Likelihood G.Squared df   p.value
## 1      -4924.091  102.2869 120 0.8772669
##
## $aggregated
##   Log.Likelihood G.Squared df   p.value
## 1      -5376.127   2.835736  3 0.4176509

br.2htm.ineq[["goodness.of.fit"]][2:3]

## $sum
##   Log.Likelihood G.Squared df   p.value
## 1      -4942.507  139.1176 120 0.1118981
##
## $aggregated
##   Log.Likelihood G.Squared df   p.value
## 1      -5376.127   2.835736  3 0.4176509
```

The results show that the 2HTM is not grossly misfitting the data as none of the likelihood ratio tests is rejected (i.e.,  $ps > .05$ ). Furthermore, as implied from the previous results, there are no differences for the aggregated data between the 2HTM with or without the order restriction applied to the guessing parameters. This, however, is not the case for the summed individual data. The log-likelihood and  $G^2$  values for the 2HTM with the order restriction is slightly worse than these values for the original 2HTM indicating that at least for some datasets the order restriction does not completely hold.

From these findings the question arises of whether or not the restrictions on the guessing parameters are justified when taking both model fit and model flexibility into account. To this end we now also consider the information criteria and again only look at the summed individual results as well as the results of the aggregated data.<sup>7</sup>

```
br.2htm$information.criteria[2:3]

## $sum
##      FIA      AIC      BIC FIA.penalty
## 1 528.8313 662.2869 2732.232   477.6878
##
## $aggregated
##      FIA      AIC      BIC FIA.penalty
## 1 26.27114 16.83574 68.58437   24.85327
```

<sup>7</sup>In previous versions of **MPTinR** the summed BIC values were reported incorrectly. Due to the non-linear nature of the  $\ln(N)$  terms in the BIC formula (Equation 9), the individual values cannot simply be summed (which is what previous versions reported). Instead, the  $G^2$  values, the  $N$ s, and the parameters, need to be summed first and BIC calculated from these summed values. From **MPTinR** version 1.9.2 onwards both `fit.mpt` et al. and `select.mpt` report correctly summed BIC values. Although FIA also contains a  $\ln(N)$  term, the FIA values *can* simply be summed (this can be easily checked by comparing the FIA of multiple individuals with a super-model in which all individuals are fit jointly with separate parameters for each individual). **MPTinR** version 1.9.2 incorrectly did not do so, but summed parts of the penalty term separately (see Equation 10) producing incorrect results. From version 1.10.2 FIA values are again correctly summed across multiple data sets.

```
br.2htm.ineq$information.criteria[2:3]

## $sum
##      FIA      AIC      BIC FIA.penalty
## 1 358.9704 699.1176 2769.063   289.4116
##
## $aggregated
##      FIA      AIC      BIC FIA.penalty
## 1 21.56424 16.83574 68.58437   20.14637
```

An alternative way of comparing models using information criteria is via function `select.mpt` (note that we call `select.mpt` with `output = "full"` to obtain the model selection table for both individual and aggregated data, the default `output = "standard"` only returns a table comparing the individual results).

```
select.mpt(list(br.2htm, br.2htm.ineq), output = "full")

##      model n.parameters G.Squared.sum df.sum      p.sum p.smaller.05
## 1      br.2htm           7      102.2869   120 0.877267           0
## 2 br.2htm.ineq           7      139.1176   120 0.111898           1
##      G.Squared.aggregated df.aggregated p.aggregated FIA.penalty.sum
## 1           2.835736           3      0.417651      477.6878
## 2           2.835736           3      0.417651      289.4116
##      delta.FIA.sum FIA.best   FIA.sum FIA.penalty.aggregated
## 1      169.8609           0 528.8313           24.85327
## 2           0.0000          40 358.9704           20.14637
##      delta.FIA.aggregated FIA.aggregated delta.AIC.sum wAIC.sum AIC.best
## 1           4.706906           26.27114      0.00000           1           40
## 2           0.000000           21.56424      36.83066           0           14
##      AIC.sum delta.AIC.aggregated wAIC.aggregated AIC.aggregated
## 1 662.2869           0           0.5           16.83574
## 2 699.1176           0           0.5           16.83574
##      delta.BIC.sum wBIC.sum BIC.best   BIC.sum delta.BIC.aggregated
## 1           0.00000           1           40 2732.232           0
## 2      36.83066           0           14 2769.063           0
##      wBIC.aggregated BIC.aggregated
## 1           0.5           68.58437
## 2           0.5           68.58437
```

The returned table compares one model per row (here split across multiple rows) using the information criteria FIA, AIC, and BIC. For each criterion, delta values (i.e., in reference to the smallest value) and absolute values are presented. The columns labeled `.best` indicate how often each model provided the best account when comparing the individual datasets. As can be seen, for FIA the model with the order restriction always provided the best account for the individual datasets (40 out of 40 individuals). In contrast, for AIC and BIC the unrestricted 2HTM only provided the best account for 26 individuals. For 14 individuals AIC and BIC were identical for both models. As the number of parameters is identical for both models and the penalty factor of AIC and BIC only includes the number of parameters as a proxy of model complexity, the difference in AIC and BIC between those two models merely reflects their differences in model fit. Furthermore, the table presents AIC and BIC weights (`wAIC` and `wBIC`; Wagenmakers and Farrell, 2004).

Overall the results indicate the utility of FIA as a measure for model selection. As expected, when the order restriction holds (as is the case for the aggregated data) FIA prefers the less complex model (i.e., the one in which the possible parameter space is reduced due to inequality restrictions). This preference for the less complex models is even evident for the cases where the order restriction might not completely hold as FIA prefers the order-restricted model for all individuals even though the order-restricted model provides a worse fit for 26 individuals. For the datasets obtained by Bröder and Schütz (2009) the more complex (i.e.,



unrestricted) 2HTM model, albeit providing the better fit, seems unjustifiably flexible when using FIA as the model selection criterion.

Note that the calculation of FIA is computationally demanding, especially when inequality restrictions are applied (see the Appendix on how the model is reparametrized for inequality restrictions). Therefore, **MPTinR** now uses the **RcppEigen** package (Bates and Eddelbuettel, 2013) for performing this step. This dramatically reduced the calculation time compared to previous versions of **MPTinR** using pure R. Furthermore, the calculation can also be split across different cores for further speed improvements. But this should only be necessary for comparatively large models.

## 4.4 Extending the Session

### 4.4.1 Model Selection on Specific Datasets

Given that for two participants (2 & 7) parameters are not uniquely identified we might want to repeat the model selection excluding those participants. To this end, `select.mpt` contains the argument `dataset` which lets us specify for which dataset we want to perform the model selection (this can be a scalar value or vector). Here we decide to perform it on all participants minus the two problematic ones:

```
select.mpt(list(br.2htm, br.2htm.ineq), output = "full", dataset = (1:40)[-c(2, 7)])
```

##	model	n.parameters	G.Squared.sum	df.sum	p.sum	p.smaller.05
## 1	br.2htm	7	98.03135	114	0.856945	0
## 2	br.2htm.ineq	7	129.08307	114	0.158253	1
##	FIA.penalty.sum	delta.FIA.sum	FIA.best	FIA.sum	delta.AIC.sum	wAIC.sum
## 1	453.8035	163.3366	0	502.8191	0.00000	1
## 2	274.9410	0.0000	38	339.4826	31.05172	0
##	AIC.best	AIC.sum	delta.BIC.sum	wBIC.sum	BIC.best	BIC.sum
## 1	38	630.0314	0.00000	1	38	2582.835
## 2	14	661.0831	31.05172	0	14	2613.887

As can be seen, this does not change the conclusions (i.e., the restricted model performs better). Furthermore, reporting of the aggregated data is suppressed now. Given that the aggregated data included all data sets (i.e., also the two now excluded ones), this makes a lot of sense as it avoids reporting of misleading results.

### 4.4.2 Considering Additional Models

As mentioned in the beginning, a popular version of the 2HTM restricts the memory parameters to be equal across item types (i.e.,  $D_o = D_n$ ). We also might be interested in whether or not this restriction is justified for our data.

```
br.2htm <- fit.mpt(d.broeder, textConnection(mod_2htm_2), list("Do = Dn"), fia = 25000)
```

```
## [1] "Computing FIA: Iteration begins at 2018-06-23 18:31:22"
```

```
## [1] "Computing FIA: Iteration stopped at 2018-06-23 18:31:23"
```

```
## Time difference of 0.3744347 secs
```

```
## Presenting the best result out of 5 minimization runs.
```

```
## [1] "Model fitting begins at 2018-06-23 18:31:23"
```

```
## Optimization routine for dataset(s) 6 7
```

```
## did not converge successfully. Tried again with use.gradient == FALSE.
```

```
## Optimization for dataset(s) 6 7
```

```
## using numerically estimated gradients produced better results. Using those results.
```

```
## Old results saved in output == 'full' [['optim.runs']].
```

```
## [1] "Model fitting stopped at 2018-06-23 18:31:26"
```

```
## Time difference of 3.66753 secs
```

```
br.2htmr.ineq <- fit.mpt(d.broeder, textConnection(mod_2htm_2),
                        list("g1 < g2 < g3 < g4 < g5", "Do = Dn"), fia = 25000)

## [1] "Computing FIA: Iteration begins at 2018-06-23 18:31:27"
## [1] "Computing FIA: Iteration stopped at 2018-06-23 18:31:50"
## Time difference of 22.81855 secs

## Presenting the best result out of 5 minimization runs.

## [1] "Model fitting begins at 2018-06-23 18:31:50"

## Optimization routine for dataset(s) 6 7
## did not converge successfully. Tried again with use.gradient == FALSE.
## Optimization for dataset(s) 6 7
## using numerically estimated gradients produced better results. Using those results.
## Old results saved in output == 'full' [['optim.runs']].

## [1] "Model fitting stopped at 2018-06-23 18:32:05"
## Time difference of 15.39986 secs
```

Given we now have four models, we again use `select.mpt` for a model selection table.

```
select.mpt(list(br.2htm, br.2htm.ineq, br.2htmr, br.2htmr.ineq), output = "full")

##          model n.parameters G.Squared.sum df.sum      p.sum p.smaller.05
## 1      br.2htm           7      102.2869   120 0.877267           0
## 2 br.2htm.ineq           7      139.1176   120 0.111898           1
## 3      br.2htmr           6      155.2634   160 0.590982           1
## 4 br.2htmr.ineq           6      182.5636   160 0.106868           2
##  G.Squared.aggregated df.aggregated p.aggregated FIA.penalty.sum
## 1          2.835736           3      0.417651          477.6878
## 2          2.835736           3      0.417651          289.4116
## 3          6.834057           4      0.144922          457.2931
## 4          6.834057           4      0.144922          276.9482
##  delta.FIA.sum FIA.best  FIA.sum FIA.penalty.aggregated
## 1    169.860909           0 528.8313          24.85327
## 2     0.000000          18 358.9704          20.14637
## 3    175.954425           0 534.9248          22.49897
## 4     9.259569          22 368.2300          17.99034
##  delta.FIA.aggregated FIA.aggregated delta.AIC.sum wAIC.sum AIC.best
## 1          4.863771          26.27114          27.02353 0.000001          11
## 2          0.156865          21.56424          63.85419 0.000000           1
## 3          4.508624          25.91599           0.00000 0.999997          29
## 4          0.000000          21.40737          27.30021 0.000001          14
##  AIC.sum delta.AIC.aggregated wAIC.aggregated AIC.aggregated
## 1 662.2869           0.000000           0.365447          16.83574
## 2 699.1176           0.000000           0.365447          16.83574
## 3 635.2634           1.998322           0.134553          18.83406
## 4 662.5636           1.998322           0.134553          18.83406
##  delta.BIC.sum wBIC.sum BIC.best  BIC.sum delta.BIC.aggregated
## 1    322.73001 0.000000           1 2732.232           5.39434
## 2    359.56067 0.000000           1 2769.063           5.39434
## 3     0.00000 0.999999          39 2409.502           0.00000
## 4    27.30021 0.000001          14 2436.802           0.00000
##  wBIC.aggregated BIC.aggregated
## 1          0.03157          68.58437
```

```
## 2      0.03157      68.58437
## 3      0.46843      63.19003
## 4      0.46843      63.19003
```

Overall the FIA results are not completely conclusive. The clear pattern is that the two models including the inequality restrictions on  $g$  are preferred compared to the ones without the inequality restriction. At the level of the individual data, the restricted 2HTM provides the best account for more participants (23/22 versus 17/18 for the unrestricted 2HTM) but its summed FIA value is somewhat smaller by around 8. On the level of the aggregated data the pattern is inverted although the difference between restricted and unrestricted 2HTM is minimal (0.17).

As a final results let us now consider the individual parameter estimates of the two best models. The parameter values of the individual data are contained in a three-dimensional array in which the first dimension corresponds to the parameter, with the second dimension we can either obtain the parameter estimates ("estimates"), the confidence intervals, or an indicator variable whether or not a given parameters is restricted or free to vary, and the third dimension corresponds to the data set.

```
str(br.2htmr$parameters$individual)

## num [1:7, 1:4, 1:40] 0.4422 0.4422 0.0294 0.1011 0.5571 ...
## - attr(*, "dimnames")=List of 3
## ..$ : chr [1:7] "Dn" "Do" "g1" "g2" ...
## ..$ : chr [1:4] "estimates" "lower.conf" "upper.conf" "restricted.parameter"
## ..$ : chr [1:40] "dataset: 1" "dataset: 2" "dataset: 3" "dataset: 4" ...
```

We can use this array to obtain mean or median estimates from the two models (i.e., selecting all parameters and all individuals, but only the parameter estimate):

```
apply(br.2htm$parameters$individual[,1,], 1, mean)

##      Dn      Do      g1      g2      g3      g4      g5
## 0.4633640 0.4730565 0.2239059 0.2927107 0.4650218 0.6549873 0.7093329

apply(br.2htm$parameters$individual[,1,], 1, median)

##      Dn      Do      g1      g2      g3      g4      g5
## 0.5151538 0.4970286 0.1337719 0.2386954 0.4870578 0.6573762 0.7618422

apply(br.2htmr$parameters$individual[,1,], 1, mean)

##      Dn      Do      g1      g2      g3      g4      g5
## 0.5085945 0.5085945 0.1592395 0.2456915 0.4226900 0.6331431 0.6919976

apply(br.2htmr$parameters$individual[,1,], 1, median)

##      Dn      Do      g1      g2      g3      g4      g5
## 0.5225297 0.5225297 0.1268384 0.2348676 0.4567806 0.6305706 0.6969520
```

These results show that  $D_o$  and  $D_n$  are quite near to each other in the unrestricted model, which suggests that the restricted model indeed provides a slightly better account. Furthermore, when restricting  $D_o$  and  $D_n$  to be equal, the other parameters do not dramatically change suggesting that this is not achieved by some sort of parameter trade-off.

Note that the mean parameter estimates are also part of the returned list:

```
br.2htm$parameters$mean

## estimates lower.conf upper.conf
```

## Dn	0.4633640	NA	NA
## Do	0.4730565	NA	NA
## g1	0.2239059	NA	NA
## g2	0.2927107	NA	NA
## g3	0.4650218	NA	NA
## g4	0.6549873	NA	NA
## g5	0.7093329	NA	NA

#### 4.4.3 Potential Problems with FIA

FIA provides an asymptotic measure of model complexity approximating normalized maximum likelihood only in the limit when the number of data points (i.e., trials for a given participant) goes to infinity. A consequence of this is that for finite sample sizes FIA may provide inaccurate or even inconsistent penalties (Navarro, 2004; Heck *et al.*, 2014). This problem is specifically prevalent when only having small numbers of trials for a given data set and cannot be overcome by collecting more participants, but only by more trials per participant. For example, imagine that in addition to the models discussed above we would also want to include the 1HT model into our set of candidate models for the data of Bröder and Schütz (2009) for which  $D_n = 0$ .

```
br.1htm <- fit.mpt(d.broeder, textConnection(mod_2htm_2), list("Dn = 0"), fia = 25000)
```

```
## Restriction starting with Dn: Constant is either equal to 0 or 1 or outside the interval
## from 0 to 1. This may lead to problems.
```

```
## Restriction starting with Dn: Constant is either equal to 0 or 1 or outside the interval
## from 0 to 1. This may lead to problems.
```

```
## [1] "Computing FIA: Iteration begins at 2018-06-23 18:32:06"
```

```
## [1] "Computing FIA: Iteration stopped at 2018-06-23 18:32:07"
```

```
## Time difference of 0.341589 secs
```

```
## Presenting the best result out of 5 minimization runs.
```

```
## [1] "Model fitting begins at 2018-06-23 18:32:07"
```

```
## [1] "Model fitting stopped at 2018-06-23 18:32:09"
```

```
## Time difference of 2.623754 secs
```

```
br.1htm.ineq <- fit.mpt(d.broeder, textConnection(mod_2htm_2),
  list("g1 < g2 < g3 < g4 < g5", "Dn = 0"), fia = 25000)
```

```
## Restriction starting with Dn: Constant is either equal to 0 or 1 or outside the interval
## from 0 to 1. This may lead to problems.
```

```
## Restriction starting with Dn: Constant is either equal to 0 or 1 or outside the interval
## from 0 to 1. This may lead to problems.
```

```
## [1] "Computing FIA: Iteration begins at 2018-06-23 18:32:10"
```

```
## [1] "Computing FIA: Iteration stopped at 2018-06-23 18:32:32"
```

```
## Time difference of 22.14903 secs
```

```
## Presenting the best result out of 5 minimization runs.
```

```
## [1] "Model fitting begins at 2018-06-23 18:32:32"
```

```
## Optimization routine for dataset(s) 6 7
```

```
## did not converge successfully. Tried again with use.gradient == FALSE.
```

```
## Optimization for dataset(s) 6 7
```

```
## using numerically estimated gradients produced better results. Using those results.
```

```
## Old results saved in output == 'full' [['optim.runs']].
```

```
## [1] "Model fitting stopped at 2018-06-23 18:32:45"
## Time difference of 13.2106 secs
```

```
select.mpt(list(br.2htm, br.2htm.ineq, br.2htmr, br.2htmr.ineq, br.1htm, br.1htm.ineq),
  output = "full")[,1:16]
```

```
##          model n.parameters G.Squared.sum df.sum    p.sum p.smaller.05
## 1      br.2htm           7      102.2869   120 0.877267           0
## 2 br.2htm.ineq           7      139.1176   120 0.111898           1
## 3      br.2htmr           6      155.2634   160 0.590982           1
## 4 br.2htmr.ineq           6      182.5636   160 0.106868           2
## 5      br.1htm           6      199.1704   160 0.019253           6
## 6 br.1htm.ineq           6      225.5127   160 0.000497           7
##  G.Squared.aggregated df.aggregated p.aggregated FIA.penalty.sum
## 1          2.835736           3      0.417651      477.6878
## 2          2.835736           3      0.417651      289.4116
## 3          6.834057           4      0.144922      457.2931
## 4          6.834057           4      0.144922      276.9482
## 5          57.389632           4      0.000000      480.4990
## 6          57.389632           4      0.000000      289.3982
##  delta.FIA.sum FIA.best  FIA.sum FIA.penalty.aggregated
## 1    169.860909           0 528.8313           24.85327
## 2     0.000000          13 358.9704           20.14637
## 3    175.954425           0 534.9248           22.49897
## 4     9.259569          22 368.2300           17.99034
## 5    221.113831           0 580.0842           23.07911
## 6     43.184129           5 402.1545           18.30159
##  delta.FIA.aggregated FIA.aggregated
## 1          4.863771          26.27114
## 2          0.156865          21.56424
## 3          4.508624          25.91599
## 4          0.000000          21.40737
## 5          30.366558          51.77393
## 6          25.589038          46.99641
```

The results reveal an interesting picture. Although the 1HTM is nested in the (unrestricted) 2HTM the FIA penalty of the summed data for the former is slightly larger than for the latter. In other words, FIA is inconsistent for those two models as the 1HTM *must* be less complex than the 2HTM. Given the number of trials per individual in the data (300 trials per individual) we cannot use FIA to compare the 2HTM and the 1HTM on the individual level. For the aggregated data this is different as we here observe the expected ordering for FIA penalties: 1HTM < 2HTM. In line with Klauer and Kellen (2011) the restricted 2HTM is somewhat less complex than the 1HTM.

To sum this up, when using FIA it is important to check if the FIA penalties are logically consistent; nested models must have lower penalties than the superordinate models. If this is not the case it may be preferable to use the aggregated data instead of the individual level data (despite the problem associated with this), as long as the parameters estimates of the aggregated data are not grossly diverging from the individual level ones. When doing so for the current data it is clear that the different 2HTM versions are clearly preferred to the 1HTM with a slight preference for the restricted 2HTM (while in all cases models with inequality restrictions on  $g$  are preferred).

## 4.5 Bootstrapping

Besides model fitting and model selection, the next major functionality of MPTinR concerns *bootstrap simulation* (Efron and Tibshirani, 1994). In the previous example using the individual data sets obtained by

Bröder and Schütz (2009), the response frequencies were low due to the small number of trials. In such cases, asymptotic statistics such as the sampling distribution of the  $G^2$  statistic or the asymptotic confidence intervals for the parameter estimates can be severely compromised (e.g., Davis-Stober, 2009). Another situation that compromises the assumptions underlying those asymptotic statistics is when parameter estimates are close to the boundaries of the parameter space (i.e., near to 0 or 1; Silvapulle and Sen, 2005). In such cases, the use of bootstrap simulations may overcome these problems.

According to the *bootstrap principle* (Efron and Tibshirani, 1994), if one assumes that an observed data sample  $\hat{F}$ , randomly drawn from a probability distribution  $F$ , provides a good characterization of the latter, then one can evaluate  $F$  by generating many random samples (with replacement) from  $\hat{F}$  and treating them as “replications” of  $\hat{F}$ . These bootstrap samples can then be used to draw inferences regarding the model used to fit the data, such as obtaining standard errors for the model’s parameter estimates ( $\hat{\Theta}$ ). When the data samples are generated on the basis of the observed data, and no assumption is made regarding the adequacy of the model to be fitted, the bootstrap is referred to as *nonparametric*. Alternatively, bootstrap samples can be based on the model’s parameters estimates that were obtained with the original data. In this case, the model is assumed to correspond to the true data-generating process and the bootstrap is designated as *parametric*. The parametric bootstraps can be used to evaluate the sampling distribution of several statistics such as the  $G^2$  and the p-values under distinct hypotheses or models (Efron and Tibshirani, 1994) (see also van de Schoot *et al.* (2010)). The use of the parametric and nonparametric bootstraps not only provides a way to overcome the limitations of asymptotic statistics, but also to evaluate parameter estimates and statistics under distinct assumptions.

MPTinR contains two higher level functions, `gen.data` and `sample.data`, that can be used for bootstrap simulations. The function `gen.data` produces bootstrap samples based on a given model and a set of parameter values. The function `sample.data` produces bootstrap samples based on a given dataset. These functions can be used separately or jointly in order to obtain parametric and nonparametric bootstrap samples. These are general-purpose functions that can be used for a wide variety of goals, such as a) obtaining confidence intervals for the estimated parameters, b) sampling distributions of the  $G^2$ -statistic and p-values under several types of null-hypotheses (van de Schoot *et al.*, 2010), and c) model-mimicry analysis (Wagenmakers *et al.*, 2004). Also, bootstrap simulations assuming individual differences, as implemented by Hu and Phillips (1999) and Moshagen (2010), can also be obtained using these functions. Both functions are calling R’s `rmultinom` function to obtain multinomially distributed random data.

Given the variety of bootstrap methods and their goals (Efron and Tibshirani, 1994), we only provide a simple example in which 200 parametric bootstrap samples are used to estimate the 95% confidence intervals for the parameter estimates obtained with the aggregated data from Bröder and Schütz (2009) (note we now only use the aggregated data but the same model again). Note, 200 is also just used for illustration purposes, in real applications values such as 1,000 or 10,000 are more appropriate.

```
br.2htm.2 <- fit.mpt(colSums(d.broeder), textConnection(mod_2htm_2))

## Presenting the best result out of 5 minimization runs.

## [1] "Model fitting begins at 2018-06-23 18:32:46"
## [1] "Model fitting stopped at 2018-06-23 18:32:46"
## Time difference of 0.07450867 secs

t(br.2htm.2[["parameters"]])

##           Dn           Do           g1           g2           g3           g4
## estimates 0.4450308 0.5561326 0.1411012 0.2152983 0.3998479 0.6045703
## lower.conf 0.3716383 0.5102111 0.1129961 0.1736149 0.3388262 0.5446874
## upper.conf 0.5184232 0.6020540 0.1692063 0.2569817 0.4608696 0.6644531
##           g5
## estimates 0.6895514
## lower.conf 0.6424590
## upper.conf 0.7366437
```

```

bs.data <- gen.data(br.2htm.2[["parameters"]][,1], 200,
                  textConnection(mod_2htm_2), data = colSums(d.broeder))

br.2htm.bs <- fit.mpt(bs.data, textConnection(mod_2htm_2), fit.agggregated = FALSE)

## Presenting the best result out of 5 minimization runs.

## [1] "Model fitting begins at 2018-06-23 18:32:47"
## [1] "Model fitting stopped at 2018-06-23 18:33:03"
## Time difference of 16.93717 secs

```

```

apply(br.2htm.bs[["parameters"]][["individual"]][,1,],
      1, quantile, probs = c(0.025, 0.975))

##           Dn           Do           g1           g2           g3           g4
## 2.5%  0.3610596  0.5074700  0.1107326  0.1793463  0.3391611  0.5471041
## 97.5% 0.5146024  0.5961463  0.1668845  0.2604401  0.4670268  0.6592800
##           g5
## 2.5%  0.6434748
## 97.5% 0.7260555

```

In this example, we first fit the original data to the (unrestricted) 2HTM to obtain parameter estimates. These estimates are displayed (along with the asymptotic confidence intervals based on the Hessian matrix) and then used as an argument to the `gen.data` function, requesting the bootstrap samples. In the next step the bootstrap samples are fitted using `fit.mpt` setting the `fit.agggregated` argument to `FALSE` to prevent MPTinR from trying to fit the (meaningless) aggregated dataset. Finally, the 95% confidence intervals are calculated by obtaining the 2.5% and 97.5% quantile from the resulting distribution of estimates for each parameter (conveniently done using R's `apply` function). As can be seen, the Hessian based confidence intervals and bootstrapped confidence intervals strongly agree indicating that the variance-covariance matrix obtained via the Hessian matrix is a good approximation of the true variance-covariance matrix (see Hu and Phillips (1999)).

Fitting the bootstrap samples takes only a few seconds. If one requests more samples or the model becomes more complicated, the multi-core functionality of MPTinR (which is more thoroughly described below and in the documentation of `fit.mpt`) can be used to reduce the fitting time by distributing the fitting across cores. For obtaining nonparametric confidence intervals the call to `gen.data` should be replaced with a call to `sample.data`. For example: `bs.data <- sample.data(d.broeder.agg, 10000, "2htm.model")`

## 4.6 Additional Functionality

This section gives an overview over the additional functions in MPTinR besides the main functions described above (see also Table 1). As already sketched above, MPTinR can fit many types of cognitive models. `fit.model` is a copy of `fit.mpt` (i.e., a model needs to be defined as strings either in a model file or via `textConnection`) with the additional possibility to specify upper and lower bounds for the parameters (as for example needed to fit SDT models). Its documentation contains examples of fitting SDT models, for example to the data of Bröder and Schütz (2009). `fit.mptinr` is a function that allows for even more flexibility in representing a model. Instead of a model as string, a model needs to be specified as a R function returning the log-likelihood of the model (known as an *objective* function) which will be minimized. This allows one to fit models to categorical data that can not be specified in model files, for example models containing integrals. The documentation of `fit.mptinr` contains an example of how to fit a SDT model to a recognition memory experiment in which the memory performance is measured via a ranking task (Kellen *et al.*, 2012). Actually, `fit.mptinr` is called by `fit.mpt` and `fit.model` with objective functions created for the models in the model file. `fit.mpt.old` is the old version of MPTinR's main function containing a different fitting algorithm (see its documentation for more information). Note that `select.mpt` accepts results from any of the fitting functions in MPTinR (as all output is produced by `fit.mptinr`). To reduce



Table 1: Overview of functions in MPTinR

Function name	Description
<code>fit.mpt</code>	fit MPT models from model files (allows computation of FIA)
<code>fit.model</code>	fit models from model files (specify parameter bounds)
<code>fit.mptinr</code>	fit models from objective function (called by <code>fit.mpt</code> )
<code>fit.mpt.old</code>	fit MPT models from model files using old fitting algorithm
<code>prediction.plot</code>	plot observed versus predicted values for fitted model
<code>select.mpt</code>	make model selection table from fitted results
<code>gen.data</code>	generate data from model and parameter values (i.e., parametric bootstrap)
<code>sample.data</code>	generate data from a given dataset (i.e., nonparametric bootstrap)
<code>gen.predictions</code>	generate response proportions from given model and parameter values
<code>make.mpt.cf</code>	returns word in $L_{\text{BMPT}}$ for MPT model from model file
<code>lbmpt.to.mpt</code>	takes word in $L_{\text{BMPT}}$ and returns MPT model file
<code>get.mpt.fia</code>	conveniently obtain FIA for an MPT model file
<code>bmpt.fia</code>	R port of <code>BMPTFIA</code> (Wu et al., 2010a)
<code>prepare.mpt.fia</code>	make string to obtain FIA in MATLAB using <code>BMPTFIA</code>
<code>make.eqn</code>	make <code>eqn</code> model file from model in <code>easy</code> format
<code>make.mdt</code>	make <code>mdt</code> data file from data <code>vector</code> , <code>matrix</code> , or <code>data.frame</code>

*Note.* Documentation containing all the arguments for each function can be obtained by typing the function name at the R prompt preceded by a `?`, e.g., `?fit.mpt`

computational time for large datasets or models, MPTinR contains the possibility to use multiple processors or a computer cluster by parallelizing the fitting algorithm using the `snowfall` package (Knaus *et al.*, 2009). Furthermore, fitting of the aggregated dataset can be disabled by setting `fit.aggregated = FALSE`.

To visualize the model fit, the results from any of the fitting functions can be passed to `prediction.plot` to visualize the model misfit. As an example we again use the fit to the aggregated data.


```
br.2htm.2 <- fit.mpt(colSums(d.broeder), textConnection(mod_2htm_2), show.messages = FALSE)

axis.labels <- c("10%-0", "90%-N", "25%-0", "75%-N", "50%-0", "50%-N",
               "75%-0", "25%-N", "90%-0", "10%-N")

prediction.plot(br.2htm.2, textConnection(mod_2htm_2), axis.labels = axis.labels,
               args.plot = list(main = "Absolute deviations (frequency scale)"))
```



### Absolute deviations (frequency scale)



This plot shows several things. First, as the trees have only two leaves (in other words, the data is binomially distributed) the misfits in one tree always cancel each other out perfectly. Second, overall, there seems to be more misfit for the old than for the new items. Third, the misfit is particularly strong for the middle category. These observations could provide inside into how to extend the model. It shows us where the model fails.

In addition to the data generating functions `gen.data` and `sample.data` described in the previous section, the function `gen.predictions` returns predicted response proportions or predicted data from a vector of parameter values for a given model. This function can be used to check if a model recovers certain parameter values (i.e., by fitting the predicted responses) and simulated identifiability (i.e., repeating this steps multiple times with random parameter values; Rouder and Batchelder, 1998). The function `gen.data` internally calls `gen.predictions`. Note that the data generating functions that take a model as an argument also work with any model that can be specified in a model file, such as signal-detection models.

`bmpt.fia` is our R port of the original `BMPTFIA` function from Wu *et al.* (2010a). As `bmpt.fia` requires a model to be entered as a word in  $L_{BMPT}$  we provide the convenience function `get.mpt.fia` which takes similar arguments as `fit.mpt` and will call `bmpt.fia` with the correct arguments. In some cases researchers might prefer to obtain the FIA from MATLAB. To this end, `MPTinR` contains the convenience function `prepare.mpt.fia`. It takes the same arguments as `get.mpt.fia` but will return a string that is the call to the original `BMPTFIA` function in MATLAB (i.e., the string just needs to be copied and pasted into MATLAB). Note however, that since using `RcppEigen` for calculating FIA the `MPTinR` implementation is faster than the corresponding Matlab code. As noted above, FIA can be directly computed in a call to `fit.mpt` (if one wants to use `select.mpt` it is necessary to use `fit.mpt` and not the other just described functions).

Finally, `MPTinR` contains several helper functions. `make.mpt.cf` will take a model file as an argument and will produce a word in  $L_{BMPT}$  using the algorithm described above. The converse can be achieved by function `lbmpt.to.mpt`<sup>8</sup> which takes a word in  $L_{BMPT}$  and produces a model file. `make.eqn` will take a model file in `easy` format and will produce a file in the EQN format. Similarly, `make.mdt` will take data (either a single `vector` or a `matrix` or `data.frame`) and will produce a single file in the MDT format containing all the datasets. EQN and MDT files are used by other programs for fitting MPT models such as `MultiTree` (Moshagen, 2010) or `HMMTree` (Stahl and Klauer, 2007).

<sup>8</sup>This function was contributed by Quentin Gronau and Franz Dietrich.

## 5 Appendix: MPTinR Algorithms

The purpose of this section is to give an overview of the algorithms used by MPTinR as they diverge from the usual employed fitting algorithm for MPT models (the EM algorithm, Hu & Batchelder, 1994). Readers mainly interested in using MPTinR for fitting MPT models may skip this section. The main task of MPTinR is model fitting, that is iteratively finding the maximum likelihood parameter estimates  $\hat{\Theta}$ . Instead of the EM algorithm, MPTinR uses the general purpose optimization algorithm implemented in R's `nlminb` function. This algorithm is a variation of Newton's method that can use the analytical or approximated (i.e., quasi-Newton) gradient or Hessian to obtain the optimal parameters within parameter bounds and is part of the PORT library (Kaufman and Gay, 2003). Previous versions of MPTinR used a different optimization algorithm (L-BFGS-B; Byrd *et al.*, 1995), which is still available in the function `fit.mpt.old`. However, as L-BFGS-B cannot use an analytical gradient we changed the algorithm to `nlminb`.

The advantages of using a general purpose optimization routine instead of a specialized one are twofold. First, the PORT routines are implemented reasonably fast in the FORTRAN programming language (compared to optimization algorithms implemented in pure R code). Second, MPTinR does not require a model to strictly follow an MPT form as described in (Riefer and Batchelder, 1988) or (Purdy and Batchelder, 2009). Instead, MPTinR literally uses the right hand sides of the model equations (e.g., Equations 1-4) and evaluates them using the current parameter values at each iteration of the optimization process. In other words, the model is not transformed into any matrix notation (see Riefer and Batchelder, 1988). Consequently, MPTinR can fit any model that can be described in a model file using (inbuilt or self-written) R functions.

When calling `fit.mpt` or `fit.model` the following steps are performed for obtaining  $\hat{\Theta}$ : The equations in the model file are parsed line by line into R **expressions** (i.e., code that can be executed). These **expressions** are concatenated to obtain the likelihood function for a given model (Equation 6). To avoid numerical underflows, the negative log of the likelihood function gives the *objective* function (i.e., the function that will be minimized). From this function the functions to calculate the gradient and the Hessian matrix of the model are derived using symbolical derivation implemented in the `D` function. The objective, gradient and Hessian function are then passed as arguments to `fit.mptinr`, the workhorse of MPTinR. `fit.mptinr` can also be called directly with an objective function if it cannot be specified in a model file.

Both the objective and the gradient function are per default passed to `nlminb` for obtaining  $\hat{\Theta}$ . In cases where `nlminb` does not converge successfully, fitting is restarted using a numerically approximated gradient (with warning). Furthermore, `fit.model` and `fit.mptinr` allows one to specify whether or not the gradient function or even the Hessian function should be passed to `nlminb` (using the Hessian function for fitting did neither improve the speed nor the quality of the fitting of MPT models and is therefore deactivated per default). To allow an assessment of the quality of the fitting algorithm, MPTinR reports the **summary** statistic of a vector containing the values of the objective function at the obtained minima for each fitting run (if `n.optim > 1`) in output element `fitting.runs`. Our experience is that the dispersion of the minima is usually 0, unless the data contains many zero cells or the model is not identified for that dataset.

When model restrictions are specified they will be applied before creating the objective function in that the expressions representing the model will be altered. In case of equality restrictions (i.e., either setting two parameters equal, e.g.,  $x = y$ , or fixing a parameter to a value e.g.,  $x = 0$ ) the model equations are altered so that the to be restricted parameter is replaced with the restriction. To apply inequality restrictions the model is altered using a variant of method A described by (Knapp and Batchelder, 2004). More specifically, each instance of the to be restricted parameter is replaced by the product of a dummy parameter and the restriction (e.g., for the restriction  $x_1 < x_2$ , each instance of  $x_1$  would be replaced with  $x_2 * \text{hank.y1}$ ). Note that all dummy parameters in MPTinR start with *hank.* and it should therefore be avoided to use parameter names starting with these characters. Whereas this reparametrization is equivalent to method A of Knapp and Batchelder, it does not preserve the  $L_{BMPT}$  structure of the model. In general it holds that for order restrictions only the rightmost element remains and all other parameters will be replaced. Note that confidence intervals around inequality restricted parameters are based on variance bounds of the parameter estimates (i.e., they represent a "worst case scenario"; Baldi and Batchelder, 2003, Equation 19).

For obtaining the FIA, MPTinR first transforms a model into a word in  $L_{BMPT}$  (if a model consists of multiple trees these are concatenated by joining parameters) which is then passed to our R port of the algorithm by Wu *et al.* (2010a). If a model is not a member of  $L_{BMPT}$  calculation of the FIA will fail. MPTinR tries to minimize computational time for the FIA by only calculating the penalty factor of the FIA

(i.e., the integral in Equation 10) as many time as needed (i.e., as many times as the ratio of the N between tree differs not as many times as N differs).

As our method of reparametrizing inequality restrictions does not preserve the MPT structure of a given model (i.e., even if it is a member of  $L_{\text{BMPT}}$  before applying the restrictions it is not thereafter), MPTinR obtains the FIA for inequality restricted models by passing the *unrestricted* model to the Wu et al. (2010a) algorithm but specifying the restrictions in the corresponding arguments. Whether or not one wants to enforce the restrictions when fitting the model is controlled by the `reparam.ineq` argument to `fit.mpt`. The default behavior is to enforce the inequality restriction by reparametrizing the model (i.e., `reparam.ineq = TRUE` is the default). This contrasts with the example given by Wu et al. in which “parameter estimates do not violate” (p. 282) the inequality restrictions and therefore the inequality restrictions are not enforced when fitting the model but the restrictions are nevertheless passed to `BMPTFIA`. This behavior can be emulated by setting `reparam.ineq = FALSE` in which case the inequality restrictions are not enforced when fitting the model, but will be taken into account when obtaining the FIA.

## References

- Akaike H (1974). “A new look at the statistical model identification.” *IEEE Transactions on Automatic Control*, **19**, 716–723. doi:10.1109/TAC.1974.1100705.
- Baldi P, Batchelder WH (2003). “Bounds on variances of estimators for multinomial processing tree models.” *Journal of Mathematical Psychology*, **47**, 467–470. doi:10.1016/S0022-2496(03)00036-1.
- Bamber D, van Santen JP (1985). “How many parameters can a model have and still be testable?” *Journal of Mathematical Psychology*, **29**(4), 443–473. doi:10.1016/0022-2496(85)90005-7.
- Batchelder WH, Riefer DM (1999). “Theoretical and empirical review of multinomial process tree modeling.” *Psychonomic Bulletin & Review*, **6**, 57–86. doi:10.3758/BF03210812.
- Bates D, Eddelbuettel D (2013). “Fast and Elegant Numerical Linear Algebra Using the RcppEigen Package.” *Journal of Statistical Software*, **52**(5), 1–24.
- Becker RA, Chambers JM, Wilks AR (1988). *The new S language: A programming environment for data analysis and graphics*. Wadsworth & Brooks/Cole Advanced Books & Software, Pacific Grove, Calif.
- Bishop YMM, Fienberg SE, Holland PW (1975). *Discrete multivariate analysis : theory and practice*. MIT Press, Cambridge, Mass.
- Bröder A, Schütz J (2009). “Recognition ROCs are curvilinear-or are they? On premature arguments against the two-high-threshold model of recognition.” *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **35**, 587. doi:10.1037/a0015279.
- Byrd RH, Lu P, Nocedal J, Zhu C (1995). “A limited memory algorithm for bound constrained optimization.” *SIAM J. Sci. Comput.*, **16**(5), 1190–1208. doi:10.1137/0916069.
- Davis-Stober CP (2009). “Analysis of multinomial models under inequality constraints: Applications to measurement theory.” *Journal of Mathematical Psychology*, **53**, 1–13. doi:10.1016/j.jmp.2008.08.003.
- Efron B, Tibshirani R (1994). *An introduction to the bootstrap*. Chapman & Hall, New York.
- Erdfelder E, Auer T, Hilbig BE, Assfalg A, Moshagen M, Nadarevic L (2009). “Multinomial Processing Tree Models.” *Zeitschrift für Psychologie / Journal of Psychology*, **217**, 108–124. doi:10.1027/0044-3409.217.3.108.
- García-Pérez MA (1994). “Parameter estimation and goodness-of-fit testing in multinomial models.” *British Journal of Mathematical and Statistical Psychology*, **47**, 247–282. doi:10.1111/j.2044-8317.1994.tb01037.x.
- Green DM, Swets JA (1966). *Signal detection theory and psychophysics*. Wiley, New York.
- Grünwald PD (2007). *The minimum description length principle*. MIT Press, Cambridge, Mass.

- Heck DW, Moshagen M, Erdfelder E (2014). “Model selection by minimum description length: Lower-bound sample sizes for the Fisher information approximation.” *Journal of Mathematical Psychology*, **60**, 29–34. ISSN 0022-2496. doi:10.1016/j.jmp.2014.06.002.
- Hu X, Batchelder WH (1994). “The statistical analysis of general processing tree models with the EM algorithm.” *Psychometrika*, **59**, 21–47. doi:10.1007/BF02294263.
- Hu X, Phillips GA (1999). “GPT.EXE: A powerful tool for the visualization and analysis of general processing tree models.” *Behavior Research Methods, Instruments, & Computers*, **31**, 220–234. doi:10.3758/BF03207714.
- Kaufman L, Gay D (2003). *The PORT Library - Optimization*. AT&T Bell Laboratories, Murray Hill, NJ. URL <http://www.bell-labs.com/project/PORT/>.
- Kellen D, Klauer KC (2011). “Evaluating models of recognition memory using first- and second-choice responses.” *Journal of Mathematical Psychology*, **55**, 251–266. doi:10.1016/j.jmp.2010.11.004.
- Kellen D, Klauer KC, Singmann H (2012). “On the Measurement of Criterion Noise in Signal Detection Theory: The Case of Recognition Memory.” *Psychological Review*, **119**, 457–479. doi:10.1037/a0027727.
- Klauer KC, Kellen D (2011). “The flexibility of models of recognition memory: An analysis by the minimum-description length principle.” *Journal of Mathematical Psychology*, **55**(6), 430–450. doi:10.1016/j.jmp.2011.09.002.
- Knapp BR, Batchelder WH (2004). “Representing parametric order constraints in multi-trial applications of multinomial processing tree models.” *Journal of Mathematical Psychology*, **48**, 215–229. doi:10.1016/j.jmp.2004.03.002.
- Knaus J, Porzelius C, Binder H, Schwarzer G (2009). “Easier Parallel Computing in R with snowfall and sfCluster.” *The R Journal*, **1**, 54–59.
- Macmillan NA, Creelman CD (2005). *Detection theory: A user’s guide*. Lawrence Erlbaum associates, New York.
- Moshagen M (2010). “multiTree: A computer program for the analysis of multinomial processing tree models.” *Behavior Research Methods*, **42**, 42–54. doi:10.3758/BRM.42.1.42.
- Muenchen RA (2012). “The Popularity of Data Analysis Software.” URL <http://r4stats.com/popularity>.
- Myung IJ, Forster MR, Browne MW (2000). “Guest Editors’ Introduction: Special Issue on Model Selection.” *Journal of Mathematical Psychology*, **44**, 1–2. doi:10.1006/jmps.1999.1273.
- Navarro DJ (2004). “A Note on the Applied Use of MDL Approximations.” *Neural Computation*, **16**(9), 1763–1768. ISSN 0899-7667. doi:10.1162/0899766041336378.
- Oberauer K (2006). “Reasoning with conditionals: A test of formal models of four theories.” *Cognitive Psychology*, **53**, 238–283. doi:10.1016/j.cogpsych.2006.04.001.
- Purdy BP (2011). “Coding graphical models: probabilistic context free languages & the minimum description length principle.” URL <http://brendanpatrickpurdy.com/resources/bppmathpsych11.pdf>.
- Purdy BP, Batchelder WH (2009). “A context-free language for binary multinomial processing tree models.” *Journal of Mathematical Psychology*, **53**, 547–561. doi:10.1016/j.jmp.2009.07.009.
- R Development Core Team (2012a). *R: A Language and Environment for Statistical Computing*. Vienna, Austria. ISBN 3-900051-07-0, URL <http://www.R-project.org/>.
- R Development Core Team (2012b). *R Data Import/Export*. 3-900051-10-0. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0. ISBN 3-900051-07-0, URL <http://cran.r-project.org/doc/manuals/R-data.html>.

- Read T, Cressie N (1988). *Goodness-of-fit statistics for discrete multivariate data*. Springer, New York.
- Riefer DM, Batchelder WH (1988). “Multinomial modeling and the measurement of cognitive processes.” *Psychological Review*, **95**, 318–339. doi:10.1037/0033-295X.95.3.318.
- Riefer DM, Batchelder WH (1991). “Statistical inference for multinomial processing tree models.” In JP Doignon, JC Falmagne (eds.), *Mathematical psychology: Current developments*, pp. 313–335. Springer, Berlin.
- Roberts S, Pashler H (2000). “How persuasive is a good fit? A comment on theory testing.” *Psychological Review*, **107**, 358–367. doi:10.1037/0033-295X.107.2.358.
- Rödiger S, Friedrichsmeier T, Kapat P, Michalke M (2012). “RKWard: A Comprehensive Graphical User Interface and Integrated Development Environment for Statistical Analysis with R.” *Journal of Statistical Software*, **49**(9), 1–34.
- Rouder JN, Batchelder WH (1998). “Multinomial models for measuring storage and retrieval processes in paired associate learning.” pp. 195–225. Lawrence Erlbaum Associates Publishers, Mahwah, NJ, US.
- Schmittmann VD, Dolan CV, Raijmakers MEJ, Batchelder WH (2010). “Parameter identification in multinomial processing tree models.” *Behavior Research Methods*, **42**, 836–846. doi:10.3758/BRM.42.3.836.
- Schwarz G (1978). “Estimating the dimension of a model.” *The annals of statistics*, **6**, 461–464. doi:10.1214/aos/1176344136.
- Silvapulle MJ, Sen PK (2005). *Constrained statistical inference: Inequality, order, and shape restrictions*. Wiley-Interscience, Hoboken, N.J.
- Snodgrass JG, Corwin J (1988). “Pragmatics of measuring recognition memory: Applications to dementia and amnesia.” *Journal of Experimental Psychology: General*, **117**(1), 34–50. doi:10.1037/0096-3445.117.1.34.
- Stahl C, Klauer KC (2007). “HMMTree: a computer program for latent-class hierarchical multinomial processing tree models.” *Behavior Research Methods*, **39**, 267–73. doi:10.3758/BF03193157.
- Stahl C, Klauer KC (2008). “A simplified conjoint recognition paradigm for the measurement of gist and verbatim memory.” *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **34**, 570–586. doi:10.1037/0278-7393.34.3.570.
- Su Y, Myung IJ, Pitt MA (2005). “Minimum Description Length and Cognitive Modeling.” In PD Grünwald, IJ Myung, MA Pitt (eds.), *Advances in minimum description length: theory and applications*, pp. 411–433. MIT Press, Cambridge, Mass.
- Valero-Mora PM, Ledesma R (2012). “Graphical User Interfaces for R.” *Journal of Statistical Software*, **49**, 1–8.
- van de Schoot R, Hoijtink H, Dekovic M (2010). “Testing Inequality Constrained Hypotheses in SEM Models.” *Structural Equation Modeling*, **17**, 443–463. doi:10.1080/10705511.2010.489010.
- Vance A (2009). “Data Analysts Captivated by R’s Power.” *The New York Times*. URL <http://www.nytimes.com/2009/01/07/technology/business-computing/07program.html>.
- Wagenmakers E, Farrell S (2004). “AIC model selection using Akaike weights.” *Psychonomic Bulletin & Review*, **11**(1), 192–196. doi:10.3758/BF03206482.
- Wagenmakers E, Ratcliff R, Gomez P, Iverson GJ (2004). “Assessing model mimicry using the parametric bootstrap.” *Journal of Mathematical Psychology*, **48**(1), 28–50. doi:10.1016/j.jmp.2003.11.004.
- Wagenmakers E, Waldorp L (2006). “Guest Editors’ Introduction: Special Issue on Model Selection.” *Journal of Mathematical Psychology*, **50**, 99–100. doi:10.1016/j.jmp.2005.01.005.

- Wu H, Myung JI, Batchelder WH (2010a). “Minimum description length model selection of multinomial processing tree models.” *Psychonomic Bulletin & Review*, **17**, 275–286. doi:10.3758/PBR.17.3.275.
- Wu H, Myung JI, Batchelder WH (2010b). “On the minimum description length complexity of multinomial processing tree models.” *Journal of Mathematical Psychology*, **54**, 291–303. doi:10.1016/j.jmp.2010.02.001.
- Zucchini W (2000). “An Introduction to Model Selection.” *Journal of Mathematical Psychology*, **44**, 41–61. doi:10.1006/jmps.1999.1276.