Selecting Amongst Multinomial Models: An Apologia for Normalized Maximum Likelihood

David Kellen

Syracuse University

Karl Christoph Klauer

Albert-Ludwigs-Universität Freiburg

Author Note

David Kellen, Department of Psychology. Karl Christoph Klauer, Department of Social Psychology and Methodology. We thank Edgar Erdfelder and Daniel W. Heck for valuable comments. Correspondence should be sent to David Kellen (davekellen@gmail.com). David Kellen received support from the Swiss National Science Foundation Grant 100014_165591. Karl Christoph Klauer was supported by DFG Reinhart-Koselleck grant DFG Kl 614/39-1.

Abstract

The modeling of multinomial data has seen tremendous progress since Riefer and Batchelder's (1988) seminal paper. One recurring challenge, however, concerns the availability of relative performance measures that strike an ideal balance between goodness of fit and functional flexibility. One approach to the problem of model selection is Normalized Maximum Likelihood (NML), a solution derived from the Minimum Description Length principle. In the present work we provide an R implementation of a Gibbs sampler that can be used to compute NML for models of joint multinomial data. We discuss the application of NML in different examples, compare NML with Bayes Factors, and show how it constitutes an important addition to researchers' toolboxes.

Keywords: Model selection, Minimum-description length, Normalized maximum likelihood, Multinomial data, Gibbs sampler

Selecting Amongst Multinomial Models: An Apologia for Normalized Maximum Likelihood

Throughout his long career, Bill Batchelder (1940-2018) championed the use of ‘cognitive measurement’ models that bridge the gap between off-the-shelf data-analytic procedures and fine-grained computational accounts. The idea is that such models can be used by both researchers and practitioners to characterize the data in terms of the main cognitive processes at play and to test substantive hypotheses. Many successful applications of this approach to modeling can be found within the class of *Multinomial Processing Tree* (MPT) models, originally proposed by Riefer and Batchelder (1988) and further developed in a series of followup papers (e.g., Hu & Batchelder, 1994; Purdy & Batchelder, 2009). MPT models assume that individuals’ categorical responses are produced by a *mixture of processes* that are assumed to occur in an *all-one-none fashion*. These processes are typically organized in a serial fashion that can be conveniently represented in a tree-like structure. Since their introduction over three decades ago, MPT models have been developed across several domains such as memory, perception, reasoning, and decision making (for reviews, see Batchelder & Riefer, 1999; Erdfelder et al., 2009; Hütter & Klauer, 2016).

As an example, let us consider a MPT model developed for the *simplified conjoint recognition paradigm* (Stahl & Klauer, 2007; see also Brainerd, Reyna, & Mojardin, 1999). In this paradigm, participants first study a list of words and subsequently engage in a recognition-memory task comprised of 1) studied targets, 2) semantically related distractors, and 3) unrelated distractors. In each recognition trial, the participant is asked to judge a single test item using one of three response alternatives: “**target**”, “**related distractor**”, and “**unrelated distractor**”. The structure of the data coming out of this paradigm is described in Table 1. Figure 1 illustrates the *simplified conjoint-recognition (SCR) model*, a simplified instantiation of the *fuzzy-trace theory* proposed by Brainerd, Reyna, and colleagues (for an overview, see Brainerd & Reyna, 2002). According to the SCR model, memory-based responses are based on the retrieval of *verbatim* and *gist traces*

Table 1
Data Structure of the Simplified Conjoint Recognition Paradigm

Item Type	Response Alternatives		
	“Target”	“Related Distractor”	“Unrelated Distractor”
Target	$\pi_{1,1}$	$\pi_{1,2}$	$\pi_{1,3}$
Related Distractor	$\pi_{2,1}$	$\pi_{2,2}$	$\pi_{2,3}$
Unrelated Distractor	$\pi_{3,1}$	$\pi_{3,2}$	$\pi_{3,3}$

Note. The π denote the probabilities associated with each item $\times$ response combination. Note that each row sums to 1.

(with probabilities denoted by V and G respectively). In the absence of any memory retrieval, a guess is produced (with guessing probabilities given by parameters a and b).

The success of measurement modeling and the MPT class in particular is sustained by the regular application of sophisticated methods coming from different statistical frameworks (e.g., Chechile, 1998; Hu & Batchelder, 1994; Klauer, 2010; Klauer & Kellen, 2011, 2015; Matzke, Dolan, Batchelder, & Wagenmakers, 2015; Smith & Batchelder, 2010; Vandekerckhove, Matzke, & Wagenmakers, 2015; Wickelmaier, & Zeileis, 2018; Wu, Myung, & Batchelder, 2010a, 2010b). One of the goals of such applications is to address a fundamental question: *which model is the best?* In some cases, the candidate models correspond to completely different theoretical accounts. In others, they state specific hypotheses regarding the processes within a given theory (e.g., which processes systematically differ across experimental conditions). The general problem of model selection is not a trivial one, given that researchers can adopt many different goals and the selection criteria adopted by researchers in their scientific practice reflect these differences (for discussions, see Kellen, 2019; Navarro, 2019). Nevertheless, we can argue that there is some level of agreement regarding a few desiderata. Specifically, a successful model should be expected to

1. *provide a good fit of the data, and*
2. *be parsimonious in the sense that:*

i its range of (highly-likely) predictions is small relative to the entire space of

Figure 1. Illustration of the SCR model proposed by Stahl and Klauer (2007). Parameters V_t and G_t denote the retrieval probabilities of verbatim and gist traces for targets. Parameters V_r and G_r denote the same retrieval probabilities for related distractors. Parameter a denotes the probability of guessing an item as either a target or a related distractor, whereas b denotes the probability that that item is subsequently guessed to be a target.

possible outcomes (i.e., a good fit is unlikely a priori), and

ii it does not postulate more processes than the ones necessary.

Many different model-selection methods have been developed with the purpose of

finding the model that strikes the best balance between these two tensions, goodness of fit and parsimony (for overviews, see the three special issues in this journal; Mulder and Wagenmakers 2016; Myung, Forster, & Browne, 2000; Wagenmakers & Waldorp 2006; see also Shiffrin, Lee, Kim, & Wagenmakers, 2008). One of these methods is *Normalized Maximum Likelihood* (NML), which stems from the Minimum Description Length framework (MDL; Grünwald, 2007; Myung, 2006; Rissanen, 1996; Wu et al., 2010a, 2010b).

The goal of the present manuscript is to advocate for the use of NML in the context of models of multinomial data such as MPTs. Using MPT models for the simplified conjoint-recognition task as examples, we will discuss some of the most popular model-selection methods in the literature and how NML has characteristics that make it an attractive option in researchers' toolboxes. We will also show how NML penalties can be conveniently approximated for any measurement model of multinomial data using an adapted Gibbs sampler, which we implemented in the statistical programming language R (R Core Team, 2019).

Some Popular Ways to Quantify Flexibility

The preference for parsimonious models leads to model-selection criteria that introduce some form of penalty (implicit or explicit) that attempts to capture the complexity of the model (Klauer, 2015). The sections below discuss some of the approaches often used by researchers to quantify flexibility. Using the SCR model as an example, we will discuss different indices in an ordering that reflects how successive indices address issues identified for the preceding ones.

Akaike and Bayesian Information Criteria

A traditional, explicit approach for penalizing models involves counting the number of free parameters and using them as an indicator of model flexibility. Two prominent examples are the *Akaike and Bayesian information criteria* (AIC and BIC; Akaike, 1973;

Schwartz, 1978). For a given model $\mathcal{M}$ with k parameters $\boldsymbol{\theta} = (\theta_1, \dots, \theta_k)$:

$$\text{AIC}_{\mathcal{M}} = \underbrace{-2 \log f_{\mathcal{M}}(\mathbf{x} \mid \hat{\boldsymbol{\theta}}_x)}_{\text{Badness of fit}} + \underbrace{2k}_{\text{Flexibility penalty}}, \quad (1)$$

$$\text{BIC}_{\mathcal{M}} = \underbrace{-2 \log f_{\mathcal{M}}(\mathbf{x} \mid \hat{\boldsymbol{\theta}}_x)}_{\text{Badness of fit}} + \underbrace{\log(n)k}_{\text{Flexibility penalty}}, \quad (2)$$

where n is the total number of observations in the data $\mathbf{x}$, and $f_{\mathcal{M}}(\mathbf{x} \mid \hat{\boldsymbol{\theta}}_x)$ is the likelihood of data $\mathbf{x}$ under the model parameters $\hat{\boldsymbol{\theta}}$ that maximize it; that is, $\hat{\boldsymbol{\theta}}_x = \text{argmax}_{\boldsymbol{\theta}} f(\mathbf{x} \mid \boldsymbol{\theta})$. In the case of both AIC and BIC, the first term quantifies the model's badness-of-fit. The second term is a penalty that attempts to quantify model flexibility using the number of free parameters as a proxy. The model with the lowest AIC/BIC is considered to be the best. One way to conveniently quantify model performance under some selection index is through weights ranging between 0 (no support) and 1 (maximum support). In the case of both AIC and BIC, the weight w for model $\mathcal{M}_i$ among J models is given by (Vandekerckhove et al., 2015):

$$w_{\mathcal{M}_i}^{(\text{IC})} = \frac{\exp((\min \text{IC}_{\mathcal{M}} - \text{IC}_{\mathcal{M}_i})/2)}{\sum_{j=1}^J \exp((\min \text{IC}_{\mathcal{M}} - \text{IC}_{\mathcal{M}_j})/2)}, \quad (3)$$

with IC denoting the specific information criterion used (e.g., AIC or BIC).

Although both criteria are being discussed together and often yield similar results (e.g., Evans, 2019; Kellen, Klauer, & Bröder, 2013), they have very distinct theoretical foundations. The AIC of a model is an estimate of the Kullback-Leibler distance between a probability distribution $f_{\mathcal{M}}$ associated with that model $\mathcal{M}$ and a distribution $f(\cdot)$ associated with the unknown, data-generating truth. The BIC difference between two models is an asymptotic approximation of their respective *Bayes Factor* under so-called unit-information priors (for overviews, see Burnham & Anderson, 2002; Vrieze, 2012).

One advantage of both AIC and BIC is the fact that they can be easily computed, which explains their widespread use. However, this ease is due to a simplified view of

flexibility, in which the number of parameters in a model is used as a proxy for how flexible it is.¹ Take the case of the SCR model illustrated in Figure 1: It postulates six free parameters that exhaust the six degrees of freedom provided by the data (see Table 1).

What this means is that, according to AIC and BIC, the model is seen as being maximally flexible. As a reference, note that one can always specify a ‘maximal’ model with as many parameters as there are degrees of freedom that can perfectly fit any observable data.²

This view of the SCR model is incorrect though. An analysis of the model’s tree structure reveals that all of its predictions comply with the following inequalities:

$$\pi_{3,1} \times \pi_{1,2} \leq \pi_{1,1} \times \pi_{3,2},$$

$$\pi_{3,2} \times \pi_{2,1} \leq \pi_{2,2} \times \pi_{3,1},$$

$$\pi_{1,3} \times \pi_{3,2} \leq \pi_{1,2} \times \pi_{3,3},$$

$$\pi_{2,3} \times \pi_{3,1} \leq \pi_{2,1} \times \pi_{3,3}.$$

To see how constrained this model is, we followed Chechile (1998) and sampled values from three Dirichlet distributions with concentration parameter $\alpha = [1, 1, 1]$. The values sampled are uniformly distributed across the probability simplexes associated with the three multinomial distributions. Only 2.8% of the sampled values complied with the inequalities stated above, indicating that the SCR model is far less complex than AIC and BIC purport it to be (after all, 100% of the samples are consistent with the maximal model).

Fisher Information Approximation

The model-selection index we will now discuss stems from the *Minimum Description Length* framework (MDL; Grünwald, 2007; Rissanen, 1996), according to which models $\mathcal{M}$

¹ More precisely, AIC and BIC rely on the number of *identifiable* parameters. If a model is non-identifiable (i.e., there is no single set of parameter values that maximizes likelihood), then it is inappropriate to count parameters. For discussions and examples, see Kellen and Klauer (in press) and Moran (2016).

² We can think of many different maximal models. As we will see later on, whether their differences are of any consequence in terms of model selection depends on the statistical framework being adopted.

and data $\mathbf{x}$ are seen as codes (e.g., strings of binary values) with some particular length L . The role of models is to provide a *compressed* representation of data by capturing the regularities present such that

$$L(\mathbf{x}) \geq L(\mathbf{x}|\mathcal{M}). \quad (4)$$

This inequality reflects the notion that any minimally competent model should be able to provide a characterization $L(\mathbf{x}|\mathcal{M})$ that is not longer than a literal description or restatement of $\mathbf{x}$. According to MDL, the best model is the one that provides the best *overall* compression by minimizing the sum

$$\mathcal{L}_{\mathcal{M}}(\mathbf{x}) = L(\mathbf{x}|\mathcal{M}) + L(\mathcal{M}). \quad (5)$$

The second term in this sum corresponds to the flexibility of the model. This criterion enforces a bias towards more parsimonious accounts. However, it would be a gross mistake to assume that this preference is due to any sort of belief regarding the true state of nature. As discussed later on, the preference for parsimonious models in the MDL framework is a byproduct of a solution to a specific optimization problem.

One of the indices coming from the MDL framework is the *Fisher Information Approximation* (FIA; Rissanen, 1996):

$$\text{FIA}_{\mathcal{M}} = \underbrace{-\log f_{\mathcal{M}}(\mathbf{x} \mid \hat{\boldsymbol{\theta}}_x)}_{\text{Badness of fit}} + \underbrace{\frac{k}{2} \log \left(\frac{n}{2\pi} \right) + \text{FIA}_f}_{\text{Flexibility penalty}}. \quad (6)$$

FIA is an asymptotic approximation to another model-selection index, Normalized Maximum Likelihood, which will be the focus of our discussion later on. Similar to BIC, FIA is comprised of two parts, one quantifying badness of fit, the other penalizing model flexibility using the total number of parameters k and sample size n . But in contrast with BIC, the flexibility penalty also includes a correction term FIA_f that captures the models' functional flexibility:

$$\text{FIA}_f = \log \int \sqrt{\det \mathcal{I}(\boldsymbol{\theta})} d\boldsymbol{\theta}, \quad (7)$$

where $\mathcal{I}(\boldsymbol{\theta})$ is the Fisher information matrix for a sample size of 1 (see Schervish, 1995).

Each element of this $k \times k$ matrix corresponds to $\mathcal{I}_{i,j}(\boldsymbol{\theta}) = \mathbb{E} \left[\frac{\partial \log f(\mathbf{x}|\boldsymbol{\theta})}{\partial \theta_i} \frac{\partial \log f(\mathbf{x}|\boldsymbol{\theta})}{\partial \theta_j} \right]$, with $\mathbb{E}[\cdot]$ denoting the expectation operator.

Let us continue with our original comparison between the SCR and the maximal model. For $n = 3 \times 100 = 300$, their respective penalties (i.e., sum of the second and third terms in Equation 6) are 13.45 and 17.11, respectively. When both models fit the data equally well, their respective weights are .98 and .02, reflecting the fact that the former model is much less flexible than the latter.³ Comparing the two models using FIA provides an alternative way to test the badness of fit of the SCR model: Because the model is defined by a set of inequalities, and no degrees of freedom remain for testing, standard frequentist methods cannot be applied (for a relevant discussion and some solutions, see Davis-Stober, 2009). Also, note that FIA can also be used to quantify model flexibility under equality (e.g., $b = .50$) and inequality constraints (e.g., $V_r \leq V_t$ and/or $G_r \leq G_t$). Resources for computing FIA are readily available for MPT models (Moshagen, 2010; Singmann & Kellen, 2013; Wu, Myung, & Batchelder, 2010a).

Although FIA overcomes the limitations associated with parameter counting by being sensitive to the flexibility coming from the models' functional forms, it still has important shortcomings. First, its computation outside of the MPT model class is far from trivial (for examples, see Klauer & Kellen, 2011), which reduces its potential user base. A second problem is related to the fact that the penalty FIA_f is an *asymptotic* penalty term, which means that it can produce quite unreasonable results under smaller sample sizes (for a discussion, see Navarro, 2004). For instance, assume that you want to compare the SCR

³ Note that FIA uses $-\log f_{\mathcal{M}}(\mathbf{x} | \hat{\boldsymbol{\theta}}_x)$ as its measure of badness of fit, in contrast with AIC/BIC, which use *two times* that same value. One consequence of this difference is that we no longer divide the relative model performances by 2 when computing model weights (see Equation 3).

model with a constrained variant assuming that there is no retrieval of verbatim traces for related distractors ($V_r = 0$). Logic implies that a constrained model is less flexible than its unconstrained version. However, FIA will sometimes state otherwise (see Heck, Moshagen, & Erdfelder, 2014; Klauer & Kellen, 2011; Navarro, 2004). FIA_f for the constrained and unconstrained SCR models are 0.53 and -1.45, respectively, which means that FIA will produce illogical penalties for any sample size n smaller than 331. Heck et al. (2014) and Klauer and Kellen (2011) report other cases in which even larger n are required. Note that this minimal sample size only guarantees that the penalties are correctly ordered for *these two models*. It says nothing about any other model comparison being conducted in tandem, especially those involving non-nested models. Also, it does not guarantee that the magnitude of the penalties provides an accurate characterization of any actual differences in model flexibility. When using FIA, researchers should evaluate its behavior across different sample sizes and check the extent to which results are dependent on the penalties used.

Bayes Factors

In Bayesian statistics, probabilities are used to express beliefs, which can be revised in light of incoming data using Bayes' Theorem (Jeffreys, 1961). Consider data $\mathbf{x}$ along with two models $\mathcal{M}_i$ and $\mathcal{M}_j$ from which they are assumed to have arisen. The posterior belief in the model given the data then is

$$P(\mathcal{M}_i | \mathbf{x}) = \frac{P(\mathbf{x} | \mathcal{M}_i)P(\mathcal{M}_i)}{P(\mathbf{x} | \mathcal{M}_i)P(\mathcal{M}_i) + P(\mathbf{x} | \mathcal{M}_j)P(\mathcal{M}_j)}. \quad (8)$$

It follows that

$$\frac{\overbrace{P(\mathcal{M}_i | \mathbf{x})}^{\text{Posterior Odds}}}{\overbrace{P(\mathcal{M}_j | \mathbf{x})}^{\text{Posterior Odds}}} = \frac{\overbrace{P(\mathbf{x} | \mathcal{M}_i)}^{\text{Bayes Factor}}}{\overbrace{P(\mathbf{x} | \mathcal{M}_j)}^{\text{Bayes Factor}}} \times \frac{\overbrace{P(\mathcal{M}_i)}^{\text{Prior Odds}}}{\overbrace{P(\mathcal{M}_j)}^{\text{Prior Odds}}}. \quad (9)$$

The Bayes Factor is a ratio of *marginal likelihoods*, which together quantify the relative support for each model that is provided by the data, with values above/below 1

indicating support for model $\mathcal{M}_i/\mathcal{M}_j$.⁴ For a model $\mathcal{M}_i$, the marginal likelihood corresponds to a weighted average of the likelihood of the data, with weights given by the parameter priors $g(\boldsymbol{\theta})$:

$$P(\mathbf{x} \mid \mathcal{M}_i) = \int f_{\mathcal{M}_i}(\mathbf{x} \mid \boldsymbol{\theta})g(\boldsymbol{\theta}) \, \mathrm{d}\boldsymbol{\theta}. \quad (10)$$

Unlike AIC/BIC or FIA, which rely on the maximum likelihood of the data under each model plus an explicit penalty term, Bayes Factors rely on *weighted average likelihoods* across the range of permissible parameter values, with weights being given by the parameter priors $g(\boldsymbol{\theta})$. This weighted averaging implicitly penalizes flexible models to the extent to which a wide range of possible data patterns are likely a priori. The model with the largest marginal likelihood is preferred.⁵ Specifically, the model likelihood and the parameter priors together yield a *predictive distribution over data*. The relative support for a model provided by some data $\mathbf{x}$ will be greater the more this predictive distribution is concentrated around $\mathbf{x}$. This distribution reflects a model's flexibility as a greater ability to accommodate data in general implies a more dispersed predictive distribution. Importantly, note that the number of parameters or their identifiability are not a concern here given that all that ultimately matters for purposes of model selection is the resulting predictive distribution of the data. Recent developments have made the computation of Bayes Factors much more accessible (e.g., Gronau, Wagenmakers, Heck, & Matzke, 2019; see also R package **TreeBUGS**; Heck, Arnold, & Arnold, 2018).

Under a Bayesian framework, the parameter priors are an integral part of a model as much as the likelihood function and thus, they play an equally important role. This is especially true in the case of Bayes Factors, given that the marginal likelihoods correspond

⁴ It is commonly asserted that the use of Bayes Factors implies a belief that one of the candidate models is the *true one*. However, this is incorrect. Bayes Factors quantify the relative change of belief amongst a set of models, regardless of their truth status (for an excellent discussion, see Morey, Romeijn, & Rouder, 2016).

⁵ Analogous to the model-selection weights described in Equation 3, the models' posterior probabilities can be used to conveniently summarize relative model performance. We will do so while assuming equal prior model probabilities $P(\mathcal{M})$ across candidate models.

Table 2
Parameter Estimates and Model-Selection Results (Stahl & Klauer, 2007, Experiment 3)

Parameter	Condition	
	1×	2×
V_t	.76 [.66 , .85]	.94 [.88 , .98]
G_t	.51 [.27 , .73]	.66 [.18 , .96]
V_r	.19 [.01 , .47]	.28 [.03 , .56]
G_r	.74 [.59 , .85]	.71 [.53 , .85]
a	.21 [.14 , .30]	
b	.18 [.13 , .23]	
Model	Bayes Factor (Post. Prob.)	ΔNML (C_{NML}^*)
Maximal	1 (.00)	0 (0)
SCR	86 (.16)	-4.05 (-4.4)
SCR _{rG}	436 (.83)	-6.03 (-6.7)
SCR _{rV}	4 (.01)	0.01 (-6.6)

Note. The reported estimates correspond to the mean and 95% quantiles of the parameter posterior distributions coming from the unconstrained SCR model. The denominator used in all Bayes Factors is the worst performing model (maximal model). Larger values indicates a greater support for the target model. Post. Prob. = Posterior model probabilities under equal prior probabilities. ΔNML is the difference between a model's NML penalty and the maximal model's. Larger ΔNML values indicate that the target model is worse. C_{NML}^* is the model's relative NML penalty, which corresponds to the penalty difference between the target model and the maximal model (because the latter is the most flexible model, the differences are always negative).

to prior-weighted average likelihoods (Lee & Vanpaemel, 2018; Liu & Aitkin, 2008). The sensitivity of Bayes Factors to the priors is especially relevant in the context of MPT modeling given the way modelers traditionally handle different tree structures. Specifically, we will show that the *customary use* of flat, so-called ‘non-informative’ priors over the [0,1] interval when working with MPT models under a Bayesian framework (e.g., Klauer, 2010; Lee & Wagenmakers, 2013; Smith & Batchelder, 2008; Vandekerckhove et al., 2015) can be quite problematic.⁶

⁶ Please keep in mind that the term ‘non-informative’ is itself controversial. On one hand, a variety of

As an example, consider the data coming from Experiment 3 reported by Stahl and Klauer (2007), in which the number of target presentations was manipulated within subjects ($1\times$ versus $2\times$). The main goal of this experiment was to test whether verbatim memory was selectively improved. The idea is that the general meaning of the word can be extracted from a single presentation, which suggests that the retrieval of gist traces should not be affected by the study repetition manipulation. In contrast, the probability of retrieving verbatim traces should increase along with the number of presentations. In the case of related distractors, the increased remembrance of what was studied should help individuals to actively dismiss them (i.e., a ‘recall-to-reject’ phenomenon; see Gallo, 2004).

We compared four models: i) a maximal, saturated model, ii) an unconstrained SCR model, and two constrained variants: iii) one in which either all gist-trace parameters were fixed to be equal across conditions (denoted SCR_{rG}), and iv) another in which all verbatim-trace parameters are restricted instead (SCR_{rV}). The maximal model adopted here characterized each three-outcome multinomial distribution m using the binary tree structure shown in Figure 2 (for similar approaches, see Kellen, Mata, & Davis-Stober, 2017; Klein, Heck, Reese, & Hilbig, 2019). Each tree m had its own set of parameters $\theta_{m,1}$ and $\theta_{m,2}$. As previously discussed, we associated each parameter with a uniform prior. The results shown in Table 2 indicate that the data support the gist-restricted SCR model over its competitors. For instance, the support for restriction of gist over verbatim traces was $435/4 \approx 108$. Also, note that these results are in line with the parameter estimates reported, where we can see clear differences in V_t . The mean differences found in V_r are consistent with what was expected, but they were generally not reliable (i.e., there was a considerable overlap between the posterior distributions of V_r).

priors (e.g., Haldane’s prior or Jeffrey’s prior) have been referred to as non-informative (e.g., Bernardo, 1979; Jeffreys, 1946). For instance, one could establish priors that are uniform across the parameter space, or instead, across the permissible range of predictions (for relevant discussions, see Heck & Wagenmakers, 2016; Segert & Davis-Stober, 2019). On the other hand, one can make the case that ‘non-informative prior’ is a misnomer in the sense all priors contain information and often end up representing beliefs that are not explicitly endorsed by the modeler (DeGroot, 1982).

Figure 2. Illustration of the maximal model for the simplified conjoint recognition task.

Now, let us consider an alternative *low-threshold model* (LT model), which postulates an item- and episodic-memory retrieval processes, with probabilities denoted by D and d , respectively. Two properties of this model are worth noticing: First, following Batchelder and Riefer (1990), it assumes that the retrieval of episodic details (i.e., recollection) is dependent on the retrieval of item memory (i.e., familiarity). In the case of targets this is captured by the probability of judging a target item based on episodic information, which corresponds to $D \times d$. Second, in line with Luce’s (1963) low-threshold model, it assumes that individuals can also retrieve supporting item information in the case of related distractors.⁷ When item information is available, accurate episodic details are also retrieved with probability d_r . The tree structure of the LT model is illustrated in Figure 3 .

The Bayes Factor comparing SCR with LT is 62.26, which indicates very little relative support for the latter model. In fact the LT model appears to perform comparably to the maximal model, with a Bayes Factor of only 1.25 relative to the maximal model. The LT’s posterior probability, which is relative to the other candidate models, is only .002. The issue with these results is that they are at odds with the fact that both SCR and LT in the present context are *equivalent* models in the sense that they have the same likelihood function (i.e., they will fit data equally well) and their parameters have a one-to-one mapping.⁸ For both targets and related distractors (i.e., ignoring subscripts):

⁷ Kellen, Erdfelder, Malmberg, Dube, and Criss (2016) discussed the suitability of a low-threshold account in paradigms like the conjoint recognition, in which related distractors are included.

⁸ The fact that both models have the same likelihood function in the context of a specific experimental

Figure 3. Illustration of the LT model for the simplified conjoint recognition task. Parameters D_t and d_t denote the retrieval probabilities of item and episodic memory retrieval for targets. Parameters D_r and d_r denote the same retrieval probabilities in the case of related distractors. Parameter a denotes the probability of guessing an item as either a target or a related distractor, whereas b denotes the probability that that item is subsequently guessed to be a target.

design does not necessarily hold across other experimental designs. One can devise experimental designs for which so-far-equivalent models make diverging predictions (e.g., Bayen, Murnane, & Erdfelder, 1996; Buchner, Erdfelder, Steffens, & Martensen, 1997; see also Schweickert & Chen, 2008). The control- and

$$\begin{aligned}
V &= Dd, \\
G &= 1 - \frac{D - Dd}{1 - Dd}, \\
D &= 1 - (1 - V)(1 - G), \\
d &= 1 - \frac{(1 - V)G}{1 - (1 - V)(1 - G)}.
\end{aligned} \tag{11}$$

The reason behind these differences is clear: The customary use of flat parameter priors yields very different prior predictive distributions. As shown in Figure 4, the observed response proportions are more similar to the majority of probability values predicted by the SCR than LT. In particular, we see that the predictive distributions of the SCR generally predict higher accuracy levels, which are in line with the observed data. However, keep in mind that there is nothing in the data themselves that is really at odds with LT. After all, the observed pattern is included in its range of predictions. It just turns out that these data are more unlikely (relative to the SCR model) under the flat priors used.

Figure 4 also illustrates the predictive distributions of the maximal model, which show a tradeoff relative to the LT model's. The responses to targets are more in line with the distribution coming from the maximal model, whereas the opposite pattern is found in the case of the related distractors. This result is interesting because it shows *how sensitive* Bayes Factors can be: The maximal model can account for *any* data in this paradigm, and the SCR/LT models can only capture a small fraction of it. Given this difference, one might expect the effect of parameter priors to be minor relative to the constraints imposed by the models' tree structure. The present results show that this belief is unwarranted.

The present example also serves to show that when relying on Bayes Factors, the way in which processes are ordered in the trees is no longer immaterial (e.g., Buchner &

automaticity-dominated models discussed later on provide an example.

Figure 4. Prior predictive distributions for targets and related distractors. Green and red squares correspond to observed proportions in the 1 $\times$ and 2 $\times$ study conditions. Note that each distribution has three categories, which means that $P(\text{“target”} \mid \cdot)$ and $P(\text{“related”} \mid \cdot)$ are sufficient to describe them.

Erdfelder, 1996; Kellen, Singmann, & Klauer, 2014).⁹ To see this, simply note that the LT model corresponds to a variant of the SCR model in which gist traces are retrieved ‘prior’

⁹ Consider the tree structure in Figure 2, which is used to represent the maximal model. The likelihood of this maximal model can be perfectly represented by many different tree structures. However, these alternative tree structures would result in different marginal likelihoods, unless one appropriately adjusts the prior distributions associated with each of them (Heck & Wagenmakers, 2016; Segert & Davis-Stober, 2019). In the case of maximal MPT models, this adjustment can be easily done using Beta priors (see Heck & Wagenmakers, 2016).

to the verbatim traces. In Figure 3, simply replace D and d with G^* and V^* , respectively in the trees for targets and related distractors. Note that the meaning of parameters differs between models: For example, parameter V_t in the SCR model represents the unconditional probability that the verbatim trace of a target is retrieved, whereas V_t^* in the variant model corresponds to the probability that the verbatim trace of a target is retrieved, *conditional* on the successful retrieval of a gist trace.

The seasoned Bayesian will in all likelihood brush off these results by pointing out that they are exactly what one should expect when computing Bayes Factors while using so-called ‘non-informative’ priors and ignoring models’ prior predictive distributions (for a related discussion, see Lee & Vanpaemel, 2018). For instance, one solution when confronted with alternative tree structures yielding the same likelihood function is to use priors that are invariant across them (Jeffrey’s prior) or to translate priors from one tree structure to the other (Heck & Wagnemakers, 2016; Segert & Davis-Stober, 2019). *We completely agree.* To be clear, we are not criticizing Bayes Factors per se but how they can lead to undesirable scenarios when researchers make use of them without taking into account their very nature. When adopting any selection criteria, researchers need to ensure that they are in line with their modeling goals and their systems of practice. Also, one should not confuse scientific goals with statistical tools nor forget that the former take precedence over the latter. In the case of Bayes Factors, its sensitivity to priors and process ordering in the trees can be at odds with researchers’ understanding of model flexibility through notions of testability and falsification that pretty much focus on the relative size of a model’s prediction space relative to the space of possible outcomes (which data are permitted and which are not; e.g., Dunn, 2000; Dunn & Anderson, 2018; Roberts & Pashler, 2000) rather than on the relative likelihood that one attaches to each possible outcome.¹⁰ Moreover, the

¹⁰ Note that the appropriateness of a prior might also depend on the experimental conditions in which they are being applied. For instance, accuracy in our example experiment was quite high, which turned out to be in line with the SCR’s predictive distribution. But had accuracy been lower, the preference for SCR (e.g., over LT) would have been lower as well.

need for specifying priors and changing them depending on the parametrization used introduces an additional requirement that can be especially difficult to satisfy when dealing with MPT models involving multiple branches and/or which are applied to data coming from a multitude of experimental conditions.

Normalized Maximum Likelihood

Our brief survey of popular model-selection indices suggests that researchers' toolboxes could benefit from having an index that *a)* successfully captures functional flexibility under any sample size, *b)* that is not sensitive to different parametrizations (e.g., process ordering), and *c)* that does not require parameter priors to be specified or adjusted. One index that achieves these desiderata is Normalized Maximum Likelihood (NML), an index coming from the MDL principle (Rissanen, 1996; Grünwald, 2007).

At the foundation of MDL are important information-theoretic results establishing the connection between code lengths and probabilities (Cover & Thomas, 2012): Let $h(\cdot)$ be a probability distribution over a set of possible outcomes X . First, the *Kraft Inequality* shows that this distribution is associated with an encoding scheme for every $\mathbf{x} \in X$ such that $L_h(\mathbf{x}) = -\log_2 h(\mathbf{x})$ bits, (rounded up to the closest integer). The larger $h(\mathbf{x})$ is, the shorter the code length. Second, *Shannon's Theorem* shows that if X is indeed distributed as $h(\cdot)$, then the encoding this distribution provides is *optimal* in the sense that the *expected* codelength will always be shorter under $h(\cdot)$ than any other distribution. A model $\mathcal{M}$ is not comprised of a single probability distribution $f_{\mathcal{M}}(\cdot)$ over X , but *a set of distributions*, each associated with a specific set of parameter values $\boldsymbol{\theta}$. When fitting a model to data $\mathbf{x} \in X$ using maximum likelihood, what we are doing is finding the distribution $f_{\mathcal{M}}(\cdot \mid \hat{\boldsymbol{\theta}}_{\mathbf{x}})$ under which data $\mathbf{x}$ are most likely, which is also the distribution in the model that compresses them the most. However, in order to determine this distribution, the data need to be known *beforehand*. This requirement indicates that such distributions are unlikely to be good candidates when attempting to determine the *single*

distribution that best compresses *incoming* data, especially when considering that this compression is likely capitalizing on some of the (non-replicable) noise found in the data; i.e., the data are being overfitted.

One conservative approach for determining the best data compression based on a model consists of finding a *universal distribution* $h_U(\cdot)$ that minimizes excess codelength under a “worst-case scenario” in which data follow a distribution $q(\cdot)$ that most compromises the ability of $h_U(\mathbf{x})$ to approximate $f_{\mathcal{M}}(\mathbf{x} \mid \hat{\boldsymbol{\theta}}_x)$ (Barron, Rissanen, & Yu, 1998; Rissanen, 2001). First, we define the *regret* function $R(h_U, \mathcal{M})$ that quantifies the difference in codelength between $h_U(\mathbf{x})$ and $f_{\mathcal{M}}(\mathbf{x} \mid \hat{\boldsymbol{\theta}}_x)$:

$$R(h_U, \mathcal{M}) = \log \left(\frac{f_{\mathcal{M}}(\mathbf{x} \mid \hat{\boldsymbol{\theta}}_x)}{h_U(\mathbf{x})} \right). \quad (12)$$

The worst-case scenario referred above is the one in which the *expected regret* $\mathbb{E}_q[R(h_U, \mathcal{M})]$ is *maximized*. The goal then is to find a universal distribution h_U that *minimizes maximum expected regret*. It can then be shown that the NML distribution h_{NML} for a given model $\mathcal{M}$ is the solution to this minimax problem (Barron et al., 1998; Rissanen, 2001):

$$h_{\text{NML}}(\cdot) = \arg \min_{h_U} \max_q \mathbb{E}_q[R(h_U, \mathcal{M})], \quad (13)$$

such that

$$h_{\text{NML}}(\mathbf{x}) = \frac{f_{\mathcal{M}}(\mathbf{x} \mid \hat{\boldsymbol{\theta}}_x)}{\sum_{\mathbf{y} \in X} f_{\mathcal{M}}(\mathbf{y} \mid \hat{\boldsymbol{\theta}}_y)}. \quad (14)$$

The probability of a given data $\mathbf{x}$ under this distribution corresponds to its maximum likelihood under $\mathcal{M}$, normalized by the sum of maximum likelihoods of all possible data $\mathbf{y} \in X$. The NML distribution for data $\mathbf{x}$ under model $\mathcal{M}$ has codelength:

$$\mathcal{L}_{\mathcal{M}}(\mathbf{x}) = -\log f_{\mathcal{M}}(\mathbf{x} \mid \hat{\boldsymbol{\theta}}_x) + \underbrace{\log \sum_{\mathbf{y} \in X} f_{\mathcal{M}}(\mathbf{y} \mid \hat{\boldsymbol{\theta}}_y)}_{\text{Flexibility penalty } C_{\text{NML}}}. \quad (15)$$

The second term C_{NML} provides a quite intuitive penalty for flexibility. A more flexible model is one with a greater capacity for successfully capturing arbitrary patterns of data (Dunn, 2000; Roberts & Pashler, 2000). One way in which this flexibility will manifest itself is in the maximum likelihoods of different data under this model, which are expected to be greater than maximum likelihoods obtained with simpler models. This difference will be directly captured in the penalty term, which will be greater for the more flexible model (note that NML will converge to the asymptotic FIA index as sample size increases). As discussed later on, it is convenient to compute the difference between the penalty of a target model and the penalty of a maximal model that always fits the data perfectly. We will denote this *relative penalty* as C_{NML}^* . This use of relative penalties does not affect model comparisons in any way. Finally, like FIA, the computation of model-performance weights can be computed for NML as described in Equation 3, with the exception that relative model performances are not divided by 2 (see Footnote 3) .

Let us once again consider the SCR and LT model. The ΔNML values reported in Table 2 replicate the qualitative pattern of results obtained with Bayes Factors, with the SCR model restricting gist traces across the different study conditions, SCR_{rG} , outperforming its competitors. Interestingly, there is a considerable degree of agreement between Bayes Factors and NML (Heck, Wagenmakers, & Morey, 2015; Shiffrin, Chandramouli, & Grünwald, 2016). In fact, NML/FIA results will converge to model selection by Bayes Factors with uninformative, parametrization-invariant, *Jeffrey's priors* (Balasubramanian, 1997). Karabatsos and Walker (2006) also discuss how NML can be cast within a Bayesian framework (but see Grünwald & Navarro, 2009).¹¹ Moreover, when comparing SCR with LT using NML, we find that they perform equally well, a result that

¹¹ Note that qualitative disagreements between Bayes Factors and NML are often associated with cases that can be deemed pathological from a scientific standpoint (for a discussion, see Shiffrin et al., 2016). But as pointed out by one reviewer, these overlaps will occur when including a maximal model among the candidates, something that is generally described as advisable, as it provides a way to evaluate badness of fit (e.g., Hilbig & Moshagen, 2014; Lee, 2016). We completely agree. However, one could circumvent this situation by establishing an alternative to the maximal model whose range of predictions is limited to the outcomes not permitted by any of the substantive candidate models (see Shiffrin et al., 2016).

reflects the fact that both models are mere reparametrizations of the same likelihood function under the present experimental design.

An important aspect of NML is that its derivation does not require any assumptions regarding the existence of a ‘true model’ nor does it call for the distributions $h_U(\cdot)$ and $q(\cdot)$ to be part of any of the candidate models. This aspect is central in the MDL framework, according to which it is unnecessary to make any ontological commitment regarding a ‘true’ data-generating distribution, when our goal is to have good ways to learn about regularities in the data and separate them from noise (Myung, Navarro, & Pitt, 2006).¹² This stance is very much in agreement with Bill Batchelder’s view of modeling in psychology, according to which the goal is not so much to reach some ultimate truth but to develop good measurement tools:

... it is time to approach cognitive modeling in simple paradigms like recognition memory from a more pragmatic, measurement-theoretic approach rather than a scientific strategy of falsifiability in service of finding the correct model. There are simply not enough degrees of freedom in the data or available behavioral measures to permit the emergence of a scientifically correct and valid model of simple recognition memory. Such efforts hit the wall a long time ago. (Batchelder & Alexander, 2013, p. 1211)

The challenge of computing NML and a Monte Carlo solution

Although NML has been applied in model-selection analyses across a number of domains (e.g., Davis-Stober & Brown, 2011; Kellen et al., 2013; Myung et al., 2006; Myung, Pitt, & Navarro, 2007; Navarro & Lee, 2005; Preacher, 2006), its popularity is somewhat limited. This situation can be attributed to the combinatorial explosion one faces when trying to compute the penalty terms. For instance, consider a multinomial distribution

¹² Despite this stance, model-recovery simulations have found to NML to select the true, data-generating model at rates that are very close to optimal (see Klauer & Kellen, 2011; Kellen, Klauer, & Bröder, 2013).

with three outcomes. With $n = 100$, there are 5151 unique response patterns that need to be considered. Now, if we have three independent distributions, as in the prototypical SCR paradigm (one per stimulus type), the number skyrockets to $5151^3 = 136,670,457,951$.

One way to sidestep this problem is to approximate the NML penalties using Monte Carlo integration methods (Evans & Swartz, 2000), as done by Klauer and Kellen (2011, 2015) when comparing different models of recognition memory (see also Roos, 2008).

Instead of a sum, the penalty term in NML can be framed as an integral over a counting measure ν that assigns measure 1 to each possible data pattern $\mathbf{y} \in X$:

$$T = \sum_{\mathbf{y} \in X} f_{\mathcal{M}}(\mathbf{y} \mid \hat{\boldsymbol{\theta}}_{\mathbf{y}}) = \int f_{\mathcal{M}}(\mathbf{y} \mid \hat{\boldsymbol{\theta}}_{\mathbf{y}}) d\nu(\mathbf{y}) \quad (16)$$

We can approximate this integral T (up to a multiplicative constant M^{-1}) by sampling data patterns $\mathbf{y}$ from a proposal distribution $g(\cdot)$ with support over X that dominates $f_{\mathcal{M}}(\cdot)$ such that there is a positive value M for which $f_{\mathcal{M}}(\mathbf{y}) \leq M \cdot g(\mathbf{y})$ for all $\mathbf{y} \in X$:

$$T_Z = \frac{1}{Z} \sum_{z=1}^Z \frac{f_{\mathcal{M}}(\mathbf{y}_z \mid \hat{\boldsymbol{\theta}}_{\mathbf{y}_z})}{Mg(\mathbf{y}_z)} \quad (17)$$

Note that Equation 17 corresponds to a standard case of *importance sampling* (see Evans & Swartz, 2000, Chap. 6). By definition, the maximum likelihood of the maximal model is always greater than the value of any probability mass function for all possible patterns.

This means that a probability mass function $g(\cdot)$ that is proportional to the maximum-likelihood function of the maximal model (whose predictions encompass all possible candidate models') would be a suitable source of proposal samples when computing NML penalties for any given set of candidate models. The question then is how to efficiently sample from a proposal density that is proportional to the maximum likelihood under the maximal model. Let $\mathbf{y} = (y_1, \dots, y_{J+1})$ be a count vector of length $J + 1$, with $n = \sum_{j=1}^{J+1} y_j$ being fixed. Note that only J categories are non-redundant as

$y_{J+1} = n - \sum_{j=1}^J y_j$. Density $g(\cdot)$ is proportional to

$$g \propto \binom{n}{y_1, \dots, y_{J+1}} \prod_{j=1}^{J+1} \left(\frac{y_j}{n} \right)^{y_j}. \quad (18)$$

When using the maximum-likelihood values under the maximal model on the right side of Equation 18 as $Mg(\mathbf{y}_z)$ in Equation 17, it follows that the multiplicative constant in the approximation T_Z is the normalizing constant that makes the maximum-likelihood values under the maximal model a probability mass function (with values summing to one across all data patterns).

The Gibbs sampler described by Klauer and Kellen (2015) allows us to sample from the proposal density $g(\cdot)$. Specifically, it takes samples for each non-redundant category j , conditional on a vector $\mathbf{y}^{(j)} = (y_1, \dots, y_{j-1}, y_{j+1}, \dots, y_J)$ comprised of all other non-redundant counts. It can be shown that each of the conditional sampling probabilities is proportional to the maximum-likelihood distribution of a maximal binomial model:

$$P(y_j | \mathbf{y}^{(j)}) \propto \binom{n^{(j)}}{y_j} \left(\frac{y_j}{n^{(j)}} \right)^{y_j} \left(\frac{n^{(j)} - y_j}{n^{(j)}} \right)^{n^{(j)} - y_j}, \quad (19)$$

where $n^{(j)} = n - \sum_{i \neq j, 1 \leq i \leq J} y_i$. As discussed in detail by Klauer and Kellen (2011), this distribution can be efficiently sampled using rejection sampling (Gelman et al., 2014). The convergence of the Gibbs sampler with regards to the density of the maximal model can be monitored by using standard methods such as the $\hat{R}$ statistic proposed by Gelman et al. (2014, p. 284).

The approximation error of an NML penalty can be obtained by computing the 99.9% confidence interval of the mean T_Z , which will include the true value with near-certainty:

$$\text{CI}_{99.9\%} = \left[T_Z - 3 \frac{s_Z}{\sqrt{Z}}, T_Z + 3 \frac{s_Z}{\sqrt{Z}} \right] \quad (20)$$

where $\frac{s_Z}{\sqrt{Z}}$ is the standard error of T_Z .¹³

A R Package for Computing NML Penalties for Multinomial Models

The Gibbs-sampling algorithm described above and the associated error-approximation estimation is implemented in the R package `NMLmulti`, which we make available on GitHub. We provide a compact example of its installation and use in the Appendix. The package can be installed directly using the following command from package `devtools`:

```
library(devtools)
install_github("davidkellen/NMLmulti")
```

The main function in this package is `run_nml`, which takes the following arguments:

- **fun**: Function with arguments `Q0` (parameters), `chd` (frequency-data vector), and `NN` (sample-size vector). Return value should be the negative log-likelihood or the G^2 statistic. Note that if the G^2 statistic is used, the penalty computed will be the *penalty difference* between the target model and the maximal model (the latter is used in computing G^2), multiplied by two relative to the likelihood-based penalties. Note that either statistics will lead to the same relative penalization of models (i.e., it won't affect your model-comparison results). However, keep in mind that this choice will determine how badness of fit is quantified (you need to use the same statistic). Importantly, the optimization routine used assumes that parameters are specified on the real scale, which means that boundaries need to imposed within the model function.¹⁴ To improve speed, the models can be previously compiled in C++ and

¹³ In practice, the Gibbs samples obtained across iterations will show dependencies, which will likely bias the estimation of s_Z . In order to address this, we estimated s_Z via *batching* (see Evans & Swartz, 2000, Chapter 7.5.2). Each batch is based on an independent chain of samples with its own starting values. The number of batches corresponds to the number of cores used in the computation. A large number of cores (at least 8-10) is advisable.

¹⁴ In our implementations, we transformed parameters into the unit interval using the distribution function of the standard normal distribution, which is given by function `pnorm()`.

included in a package using functions from `Rcpp` and `devtools` (see Appendix).

- **par1**: Number of free parameters in the model.
- **ks**: Vector, with each element indicating the number of categories in a multinomial distribution. Note that the length of this vector will indicate the number of multinomial distributions.
- **Ns**: Vector, with each element indicating the sample size n of each multinomial distribution. Needs to have the same length as **ks**.
- **fits**: Number of fit runs per dataset (only best fit is retained). Default is 2. NML requires the determination of the maximum likelihood for each dataset considered. Therefore, you should set the number of fits to a number that is very likely to result in the best possible fit (i.e., that the global maximum is found).
- **batchsize**: Number of samples used (per core) in each update of T_Z . Default is 1000.
- **thin**: Number of consecutive samples discarded in a sequence. Default is 1. Note that **batchsize** is preserved, which means that by increasing the thinning, one is increasing the total number of data samples taken. For example, if **batchsize** = 1000 and **thin** = 5, then 5000 samples are taken (per core), but only 1000 are ultimately retained for fitting.
- **burn**: Scalar, number of initial samples to be taken and discarded before sampling the first batch. For each multinomial distribution, the initial frequency samples are taken from a noninformative Dirichlet-multinomial distribution with parameters $\alpha_j = 1$.
- **cores**: Scalar, number of cores to be used in the computation of T_Z . Default is `NULL`, which corresponds to 80% of the cores available.
- **precision**: Scalar, length of the 99.9% confidence interval of $\log(T_Z)$. Default value is 0.20. This value is deemed reasonable given that any penalized-fit differences

below 0.5 are pretty much irrelevant for purposes of model selection (e.g., their effect in terms of model-selection weights is negligible).

The function will run until the pre-specified level of precision is reached, and return *a)* the NML estimate, *b)* the precision, and *c)* the number of total samples. For example, the NML penalty for the unconstrained SCR model applied to Stahl and Klauer’s (2007) Experiment 3 can be obtained as follows:

```
nml_scr <- run_nml(fun=examplemodels::SCR,
                  parl=10,
                  ks=rep(3,5),Ns=c(rep(100,4),200), fits=1,
                  cores=6,batchsize=5000,burn=10000,precision=0.1)
```

Although not necessary, note that the model here is set as function inside a R package entitled `examplemodels` (for a discussion, see below as well as the Appendix). When using a precision of 0.10, the penalty estimates for SCR/LT converged in around ten minutes when using a desktop computer equipped with a twelve-core Intel® Xeon® processor with 3.07 GHz clock speed (using six cores). Each restricted version of SCR took around 20 minutes. This difference in computation time is due to the fact that the proposal density $g(\cdot)$ used is more efficient the more flexible a model is (as its likelihood function becomes increasingly similar to the maximal model’s). Also, the models were implemented in C++ and linked to R via `Rcpp` (Eddelbuettel et al., 2011). Although not required, this leads to non-negligible improvements in terms of computation time. We discuss this implementation in the Appendix. Finally, as a sanity check, we compared our estimates with penalties computed exactly in previous studies. For instance, Klauer and Kellen (2011) computed exact penalties for the one- and two-high threshold models in the context of a second-choice recognition paradigm. Across conditions, the penalty differences between both models were -0.69 and -1.12. The respective estimated differences were -0.70 and -1.17. These results indicate that our algorithm provides a good approximation of the penalties.

NML Application: Selecting Amongst MPT Models of Control and Automaticity in Weapon Misidentification

In the previous example, we focused on different parameterizations of a saturated model of recognition memory. In our next application, we will compare three MPT models designed for process-dissociation designs (Buchner, Erdfelder, & Vaterrodt-Plünnecke, 1993; Jacoby, 1991). In this particular case, we will focus on a process-dissociation study concerning the role of automatic and controlled processes in the identification of objects as weapons (Bishara & Payne, 2009). Process-dissociation designs are used to disentangle the contribution of (presumably-independent) processes assumed to underlie a individuals' responses. This is achieved by contrasting *congruent* trials in which the different processes are assumed to support the same response, with *incongruent* trials in which they are expected to support opposite responses (for discussions, see Buchner & Erdfelder, 1996; Graf & Komatsu, 1994; Klauer, Dittrich, Scholtes, & Voss, 2015; see also Footnote 8). Consider the study by Payne, Lambert, and Jacoby (2002): At each trial, participants were shown a target item, which could be a tool or a weapon. Prior to the target presentation, participants were primed with a face that was either white or black. Payne et al. also varied the response deadline across trials (200ms, 450ms, and 700ms), which is expected to selectively influence the controlled processes (see Payne, 2001; Yonelinas & Jacoby, 1994).

The data reported by Payne et al. (2002) was used by Bishara and Payne (2009) to compare three different MPT models that instantiate alternative views about the relationship between automatic and controlled processes.¹⁵ The first two models considered assume two processes only, one controlled and one automatic, whose probabilities of occurring are denoted by C and A , respectively. The *control-dominating model* (CDM) assumes that automatic processes can only drive responses when the control process is absent (see Figure 5, upper panel). In contrast, the *automaticity-dominant model* (ADM)

¹⁵ Bishara and Payne (2009) included other studies. However, the one originally conducted by Payne et al. (2002) stood out as particularly interesting given that it led to large differences in badness of fit that were subsequently 'reversed' when using AIC/BIC.

Figure 5. Illustration of CDM and ADM. Note that unlike the previous examples where we specified a tree per stimulus type, here we represent each model by a single tree and describe the response mappings across stimulus class in the terminal nodes. Parameters C and A denote the probabilities that the controlled and automatic processes take place. Parameter B denotes the probability of guessing the target stimulus to be a tool. W = White, and B = Black. T = Tool, and G = Gun. The + and - signs indicate response accuracy (correct and incorrect, respectively) at each prime-target combination.

postulates that control processes can only intervene in the absence of an automatic process (see Figure 5, lower panel). In both models, when neither process plays a role, a guessing-based response is expected.

The third MPT model considered was the *Quadruple (Quad) Process Model* originally proposed by Conrey, Sherman, Gawronski, Hugenberg, and Groom (2005). According to the Quad-model, which is illustrated in Figure 6, the presentation of the prime leads to the

Figure 6. Illustration of the Quad-Model. Parameters C and A denote the probabilities that the controlled and automatic processes take place. Parameter B denotes the probability of guessing the target stimulus to be a tool. W = White, and B = Black. T = Tool, and G = Gun. The + and - signs indicate response accuracy (correct and incorrect, respectively) at each prime-target combination.

activation of a stereotypical association with probability S .¹⁶ Independently of whether this activation takes place, the true status of the target item is determined with probability D . In some cases, this determination conflicts with the activated association: With probability OB , the latter process overcomes the bias coming from the former. With probability $1 - OB$, it does not. The association process is assumed to be automatic, whereas the target-determination and overcoming-bias processes are assumed to be controlled. Similar to the previous models, the Quad-model also assumes that in the absence of any of these processes, a guessing-based response is given.

¹⁶ Parameter S is normally denoted by AC . We use S instead in order to avoid any confusion with the A and C parameters of the CDM and ADM.

Bishara and Payne (2009) compared the performance of the three models using AIC and BIC. Although the Quad-model always matched or outperformed the CDM and ADM in terms of badness of fit, both AIC and BIC almost invariably found it to be the worst-performing of the three. In general, the ADM was found to be the superior model. Results are reported in Table 3. The reason for the Quad model’s weak performance is that it has three additional parameters, which leads to much more severe AIC/BIC penalties (the penalty differences were 6 and 17.59, respectively).¹⁷ Overall, these results speak strongly against the Quad-model, which is judged to be an overly flexible model. To see how extreme this difference is, note that in cases where the Quad-model fits the data as well as the ADM/CDM, its weight will be virtually zero. The only cases in which it could succeed are those in which it fits the data well whereas the ADM/CDM grossly misfit them.

As previously discussed, there are good reasons to generally question the use of model-selection indices that rely on the number of parameters as a proxy for flexibility. In this particular case however, additional concerns arise when investigating the tree structure of the Quad-model. As discussed by Bishara and Payne (2009), the Quad-Model can be recast as a mixture of the CDM and ADM with weights OB and $1 - OB$. This insight explains why the Quad-model always matched or outperformed both ADM and CDM in terms of badness of fit. However, this insight also reveals that the Quad-model predictions cannot go beyond a mixture of the CDM and ADM predictions, which are in many respects quite similar. Despite the large difference in the number of parameters, it seems likely that the models do not differ so much in terms of their flexibility. We computed each model’s respective NML penalty under Payne et al.’s (2002) experimental design and found our suspicion to be corroborated. As shown in Table 3, the Quad-model is only slightly more flexible than the ADM (under equal fit, their weights would be .35 and .65, respectively), which in turn is also slightly more complex than the CDM (.38 and .62; note

¹⁷ Because some trials were excluded by Payne et al. (2002), we report the BIC penalty for the average number of trials per participant.

Table 3

Model performance: Summed selection indices, percentage of individual ‘wins’, average model-selection weights w , for AIC, BIC, and NML, as well as the relative NML penalties C_{NML}^ , for Payne et al.’s (2002) data*

Model	AIC	$w^{(AIC)}$	BIC	$w^{(BIC)}$	NML	$w^{(NML)}$	C_{NML}^*
ADM	1881 (62%)	.49	4012 (64%)	.53	-605 (13%)	.26	-10.8
CDM	1916 (35%)	.43	4047 (36%)	.47	-633 (57%)	.37	-11.3
Quad	2244 (3%)	.08	5442 (0%)	.00	-643 (30%)	.37	-10.2

that the the CDM and ADM are deemed equally flexible according to AIC/BIC). These differences completely overturn the previous model-selection results. First, the ADM no longer outperforms the CDM. Second, the Quad-model is no longer dismissed: It wins in terms of summed NML and is the best-performing model for a non-negligible number of individual datasets.

Discussion

The goal of the present work is to make a case for the use of NML when selecting among models of multinomial data, which includes the MPT model class originally proposed by Riefer and Batchelder (1988). By quantifying model flexibility in terms of the general ability of a model to fit any possible data under a given experimental design, NML overcomes the limitations of parameter-counting approaches such as AIC and BIC, without requiring researchers to make the kind of prior commitments that Bayes Factors require. Moreover, the derivation of NML explicitly emphasizes models’ ability to provide a good compression of the data, irrespective of their truth status. These characteristics make NML a quite attractive option: It relies on a quite intuitive notion of model flexibility, and its clear lack of ontological commitments highlights the fact that none of our models is more than a (hopefully useful) approximation. The examples we covered here demonstrate how NML changes the way that one usually thinks about models and their relation with data. In the case of Stahl and Klauer’s (2007) data, it showed that even the most flexible SCR model which assumes separate parameters in each experimental condition is quite

constrained. It also captured the fact that an apparently distinct LT model is in fact nothing more than a reparametrization. Moreover, the reanalysis of Payne et al.’s (2002) data showed how NML can overturn previous model-selection results that hinged on penalties based on parameter counting. NML comparisons in all of these examples were implemented using the R package `NMLmulti`, which we make available here.

Although NML is a well-established index, it is still the target of many research efforts, with different refinements and extensions being proposed (for a recent survey, see Grünwald & Roos, 2019). Among recent developments, one of the most promising ones concerns the introduction of a *data prior* (e.g., Zhang, 2012), which among other things can be used to constrain data to a region of permissible values (e.g., Dunn & Anderson, 2018). For example, one could constrain data to be at or above chance, or to conform to a certain number of axioms (e.g., transitivity; Cavagnaro & Davis-Stober, 2014). One of the advantages of introducing such constraints is that it would enable the application of NML to model comparisons involving unbounded random variables (see Zhang, 2012). It should be noted that the MPT model family has been recently extended in order to accommodate other types of data such as mouse-tracking data and reaction times (Heck, Erdfelder, & Kieslich, 2018; Klauer & Kellen, 2018).

Finally, despite the attractive features of NML and the MDL framework in general, it is important to keep in mind that it is only one of many solutions currently available in researchers’ toolboxes. Which model-selection index one should adopt ultimately depends on a set of general evaluation criteria, which can vary according to the researchers’ goals (Kellen, 2019; Navarro, 2019). Also, there is the matter of feasibility: For instance, unlike Bayes Factors, NML cannot be realistically applied in the context of hierarchical models. No matter how sophisticated and/or virtuous any method might be, it will never excuse the researcher of reasoning about their target research problem and the numerous constraints they face. Not surprisingly, we find this closing thought echoing earlier words by Batchelder and colleagues:

model selection lies at the core of the scientific inference process. Accordingly, a theoretically well-justified and widely applicable methodology can help advance science. We believe that MDL represents such a methodology and that it provides versatile yet powerful tools for assessing the validity of MPT models [...] Statistical model-selection techniques alone, however sophisticated, are not a panacea for all inference problems. Other, nonstatistical means of model evaluation, such as plausibility, interpretability, and explanatory adequacy, are equally—if not more— important. Instead of automatic tools implemented in software, statistical model-selection methods can be most useful if they are combined with the judicious use of sound subjective but scientific judgment. Wu, Myung, & Batchelder (2010a, p. 284)

References

- Akaike, H. (1973). Information theory as an extension of the maximum likelihood principle. In B. N. Petrov & F. Csaki (Eds.), *Second international symposium on information theory* (p. 267-281). Budapest: Akademiai Kiado.
- Balasubramanian, V. (1997). Statistical inference, occam's razor, and statistical mechanics on the space of probability distributions. *Neural computation*, 9, 349–368.
- Barron, A., Rissanen, J., & Yu, B. (1998). The minimum description length principle in coding and modeling. *IEEE Transactions on Information Theory*, 44, 2743–2760.
- Batchelder, W. H., & Alexander, G. E. (2013). Discrete-state models: comment on pazzaglia, dube, and rotello (2013). *Psychological Bulletin*, 139, 1204-1212.
- Batchelder, W. H., & Riefer, D. M. (1999). Theoretical and empirical review of multinomial process tree modeling. *Psychonomic Bulletin & Review*, 6, 57–86.
- Bayen, U. J., Murnane, K., & Erdfelder, E. (1996). Source discrimination, item detection, and multinomial models of source monitoring. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 197-215.
- Bernardo, J. M. (1979). Reference posterior distributions for bayesian inference. *Journal of the Royal Statistical Society: Series B (Methodological)*, 41, 113–128.
- Bishara, A. J., & Payne, B. K. (2009). Multinomial process tree models of control and automaticity in weapon misidentification. *Journal of Experimental Social Psychology*, 45, 524–534.
- Brainerd, C. J., & Reyna, V. F. (2002). Fuzzy-trace theory and false memory. *Current Directions in Psychological Science*, 11, 164–169.
- Brainerd, C. J., Reyna, V. F., & Mojardin, A. (1999). Conjoint recognition. *Psychological review*, 106, 160-179.
- Buchner, A., & Erdfelder, E. (1996). On assumptions of, relations between, and evaluations of some process dissociation measurement models. *Consciousness and cognition*, 5, 581–594.

- Buchner, A., Erdfelder, E., Steffens, M. C., & Martensen, H. (1997). The nature of memory processes underlying recognition judgments in the process dissociation procedure. *Memory & Cognition*, *25*, 508–517.
- Buchner, A., Erdfelder, E., & Vaterrodt-Plünnecke, B. (1995). Toward unbiased measurement of conscious and unconscious memory processes within the process dissociation framework. *Journal of Experimental Psychology: General*, *124*, 137–160.
- Burnham, K. P., & Anderson, D. R. (2002). *Model selection and multimodel inference: a practical information-theoretic approach*. New York, NY: Springer.
- Cavagnaro, D. R., & Davis-Stober, C. P. (2014). Transitive in our preferences, but transitive in different ways: An analysis of choice variability. *Decision*, *1*, 102–122.
- Chechile, R. A. (1998). A new method for estimating model parameters for multinomial data. *Journal of Mathematical Psychology*, *42*, 432–471.
- Conrey, F. R., Sherman, J. W., Gawronski, B., Hugenberg, K., & Groom, C. J. (2005). Separating multiple processes in implicit social cognition: the quad model of implicit task performance. *Journal of personality and social psychology*, *89*, 469–487.
- Cover, T. M., & Thomas, J. A. (2012). *Elements of information theory*. John Wiley & Sons.
- Davis-Stober, C. P., & Brown, N. (2011). A shift in strategy or "error"? strategy classification over multiple stochastic specifications. *Judgment and Decision Making*, *6*, 800–813.
- DeGroot, M. H. (1982). Lindley's paradox: Comment. *Journal of the American Statistical Association*, *77*, 336–339.
- Dunn, J. C. (2000). Model complexity: The fit to random data reconsidered. *Psychological Research*, *63*, 174–182.
- Dunn, J. C., & Anderson, L. (2018). Signed difference analysis: Testing for structure under monotonicity. *Journal of Mathematical Psychology*, *85*, 36–54.
- Eddelbuettel, D., François, R., Allaire, J., Ushey, K., Kou, Q., Russel, N., . . . Bates, D.

- (2011). Rcpp: Seamless r and c++ integration. *Journal of Statistical Software*, 40, 1–18.
- Erdfelder, E., Auer, T.-S., Hilbig, B. E., Aßfalg, A., Moshagen, M., & Nadarevic, L. (2009). Multinomial processing tree models: A review of the literature. *Zeitschrift für Psychologie/Journal of Psychology*, 217, 108–124.
- Evans, M., & Swartz, T. (2000). *Approximating integrals via monte carlo and deterministic methods*. Oxford, UK: Oxford University Press.
- Evans, N. J. (2019). Assessing the practical differences between model selection methods in inferences about choice response time tasks. *Psychonomic bulletin & review*, 1–29.
- Gallo, D. A. (2004). Using recall to reduce false recognition: diagnostic and disqualifying monitoring. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30, 120–128.
- Graf, P., & Komatsu, S. (1994). Process dissociation procedure: Handle with caution! *European Journal of Cognitive Psychology*, 6, 113–129.
- Gronau, Q. F., Wagenmakers, E.-J., Heck, D. W., & Matzke, D. (2019). A simple method for comparing complex models: Bayesian model comparison for hierarchical multinomial processing tree models using warp-iii bridge sampling. *Psychometrika*, 84, 261–284.
- Grünwald, P., & Navarro, D. J. (2009). Nml, bayes and true distributions: A comment on karabatsos and walker (2006). *Journal of Mathematical Psychology*, 53, 43–51.
- Grünwald, P., & Roos, T. (2019). Minimum Description Length Revisited. *arXiv e-prints*, arXiv:1908.08484.
- Grünwald, P. D. (2007). *The minimum description length principle*. Boston, MA: MIT press.
- Heck, D. W., Arnold, N. R., & Arnold, D. (2018). Treebugs: An r package for hierarchical multinomial-processing-tree modeling. *Behavior Research Methods*, 50, 264–284.
- Heck, D. W., Erdfelder, E., & Kieslich, P. J. (2018). Generalized processing tree models:

- Jointly modeling discrete and continuous variables. *Psychometrika*, 83, 893–918.
- Heck, D. W., 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.
- Heck, D. W., & Wagenmakers, E.-J. (2016). Adjusted priors for bayes factors involving reparameterized order constraints. *Journal of Mathematical Psychology*, 73, 110–116.
- Heck, D. W., Wagenmakers, E.-J., & Morey, R. D. (2015). Testing order constraints: Qualitative differences between bayes factors and normalized maximum likelihood. *Statistics & Probability Letters*, 105, 157–162.
- Hu, X., & Batchelder, W. H. (1994). The statistical analysis of general processing tree models with the em algorithm. *Psychometrika*, 59, 21–47.
- Hütter, M., & Klauer, K. C. (2016). Applying processing trees in social psychology. *European Review of Social Psychology*, 27, 116–159.
- Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of memory and language*, 30, 513–541.
- Jeffreys, H. (1946). An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 186, 453–461.
- Jeffreys, H. (1998). *The theory of probability*. Oxford, UK: Oxford University Press.
- Karabatsos, G., & Walker, S. G. (2006). On the normalized maximum likelihood and bayesian decision theory. *Journal of Mathematical Psychology*, 50, 517–520.
- Kellen, D. (2019). A model hierarchy for psychological science. *Computational Brain & Behavior*, 2, 160–165.
- Kellen, D., Erdfelder, E., Malmberg, K. J., Dubé, C., & Criss, A. H. (2016). The ignored alternative: An application of luce’s low-threshold model to recognition memory. *Journal of Mathematical Psychology*, 75, 86–95.
- Kellen, D., & Klauer, K. C. (2015). Signal detection and threshold modeling of

- confidence-rating rocs: A critical test with minimal assumptions. *Psychological Review*, 122, 542.
- Kellen, D., & Klauer, K. C. (in press). Theories of the wason selection task: A critical assessment of boundaries and benchmarks. *Computational Brain & Behavior*.
- Kellen, D., Klauer, K. C., & Bröder, A. (2013). Recognition memory models and binary-response rocs: A comparison by minimum description length. *Psychonomic Bulletin & Review*, 20, 693–719.
- Kellen, D., Mata, R., & Davis-Stober, C. P. (2017). Individual classification of strong risk attitudes: An application across lottery types and age groups. *Psychonomic bulletin & review*, 24, 1341–1349.
- Kellen, D., Singmann, H., & Klauer, K. C. (2014). Modeling source-memory overdistribution. *Journal of Memory and Language*, 76, 216–236.
- Klauer, K. C. (2010). Hierarchical multinomial processing tree models: A latent-trait approach. *Psychometrika*, 75, 70–98.
- Klauer, K. C. (2015). Mathematical modeling. In B. Gawronski & G. V. Bodenhausen (Eds.), *Theory and explanation in social psychology* (p. 371-389). New York, NY: Guilford Press.
- Klauer, K. C., Dittrich, K., Scholtes, C., & Voss, A. (2015). The invariance assumption in process-dissociation models: An evaluation across three domains. *Journal of Experimental Psychology: General*, 144, 198-221.
- Klauer, K. C., & 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.
- Klauer, K. C., & Kellen, D. (2015). The flexibility of models of recognition memory: The case of confidence ratings. *Journal of Mathematical Psychology*, 67, 8–25.
- Klauer, K. C., & Kellen, D. (2018). Rt-mpts: Process models for response-time distributions based on multinomial processing trees with applications to recognition

- memory. *Journal of Mathematical Psychology*, 82, 111 - 130.
- Klein, S. A., Heck, D. W., Reese, G., & Hilbig, B. E. (2019). On the relationship between openness to experience, political orientation, and pro-environmental behavior. *Personality and Individual Differences*, 138, 344–348.
- Lee, M. D., & Vanpaemel, W. (2018). Determining informative priors for cognitive models. *Psychonomic Bulletin & Review*, 25, 114–127.
- Liu, C. C., & Aitkin, M. (2008). Bayes factors: Prior sensitivity and model generalizability. *Journal of Mathematical Psychology*, 52, 362–375.
- Luce, R. D. (1963). A threshold theory for simple detection experiments. *Psychological Review*, 70, 61-79.
- Matzke, D., Dolan, C. V., Batchelder, W. H., & Wagenmakers, E.-J. (2015). Bayesian estimation of multinomial processing tree models with heterogeneity in participants and items. *Psychometrika*, 80, 205–235.
- Moran, R. (2016). Thou shalt identify! the identifiability of two high-threshold models in confidence-rating recognition (and super-recognition) paradigms. *Journal of Mathematical Psychology*, 73, 1–11.
- Morey, R. D., Romeijn, J.-W., & Rouder, J. N. (2016). The philosophy of bayes factors and the quantification of statistical evidence. *Journal of Mathematical Psychology*, 72, 6 - 18.
- Moshagen, M. (2010). multitree: A computer program for the analysis of multinomial processing tree models. *Behavior Research Methods*, 42, 42–54.
- Mulder, J., & Wagenmakers, E.-J. (2016). Editors’ introduction to the special issue “bayes factors for testing hypotheses in psychological research: Practical relevance and new developments”. *Journal of Mathematical Psychology*, 72, 1–5.
- Myung, I. J., Forster, M. R., & Browne, M. W. (2000). Model selection [special issue]. *Journal of Mathematical Psychology*, 44, 37.
- Myung, J. I., Navarro, D. J., & Pitt, M. A. (2006). Model selection by normalized

- maximum likelihood. *Journal of Mathematical Psychology*, 50, 167–179.
- Myung, J. I., Pitt, M. A., & Navarro, D. J. (2007). Does response scaling cause the generalized context model to mimic a prototype model? *Psychonomic Bulletin & Review*, 14, 1043–1050.
- Navarro, D. J. (2004). A note on the applied use of mdl approximations. *Neural Computation*, 16, 1763–1768.
- Navarro, D. J. (2019). Between the devil and the deep blue sea: Tensions between scientific judgement and statistical model selection. *Computational Brain & Behavior*, 2, 28–34.
- Navarro, D. J., & Lee, M. D. (2005). An application of minimum description length clustering to partitioning learning curves. In *Proceedings of the international symposium on information theory* (pp. 587–591).
- Payne, B. K. (2001). Prejudice and perception: The role of automatic and controlled processes in misperceiving a weapon. *Journal of personality and social psychology*, 81, 181–192.
- Payne, B. K., Lambert, A. J., & Jacoby, L. L. (2002). Best laid plans: Effects of goals on accessibility bias and cognitive control in race-based misperceptions of weapons. *Journal of Experimental Social Psychology*, 38, 384–396.
- Preacher, K. J. (2006). Quantifying parsimony in structural equation modeling. *Multivariate Behavioral Research*, 41, 227–259.
- Purdy, B. P., & Batchelder, W. H. (2009). A context-free language for binary multinomial processing tree models. *Journal of Mathematical Psychology*, 53, 547–561.
- R Core Team. (2019). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from <http://www.R-project.org/>
- Riefer, D. M., & Batchelder, W. H. (1988). Multinomial modeling and the measurement of cognitive processes. *Psychological Review*, 95, 318–339.
- Rissanen, J. (2001). Strong optimality of the normalized ml models as universal codes and

- information in data. *IEEE Transactions on Information Theory*, 47, 1712–1717.
- Rissanen, J. J. (1996). Fisher information and stochastic complexity. *IEEE transactions on information theory*, 42, 40–47.
- Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? a comment on theory testing. *Psychological review*, 107, 358–367.
- Roos, T. (2008). Monte carlo estimation of minimax regret with an application to mdl model selection. In *2008 ieee information theory workshop* (pp. 284–288).
- Schervish, M. J. (2012). *Theory of statistics*. New York, NY: Springer.
- Schwarz, G. (1978). Estimating the dimension of a model. *The annals of statistics*, 6, 461–464.
- Schweickert, R., & Chen, S. (2008). Tree inference with factors selectively influencing processes in a processing tree. *Journal of Mathematical Psychology*, 52(3), 158–183.
- Segert, S., & Davis-Stober, C. P. (2019). A general approach to prior transformation. *Journal of Mathematical Psychology*, 91, 103–118.
- Shiffrin, R. M., Chandramouli, S. H., & Grünwald, P. D. (2016). Bayes factors, relations to minimum description length, and overlapping model classes. *Journal of mathematical psychology*, 72, 56–77.
- Shiffrin, R. M., Lee, M. D., Kim, W., & Wagenmakers, E.-J. (2008). A survey of model evaluation approaches with a tutorial on hierarchical bayesian methods. *Cognitive Science*, 32, 1248–1284.
- Singmann, H., & Kellen, D. (2013). Mptinr: Analysis of multinomial processing tree models in r. *Behavior Research Methods*, 45, 560–575.
- Smith, J. B., & Batchelder, W. H. (2010). Beta-mpt: Multinomial processing tree models for addressing individual differences. *Journal of Mathematical Psychology*, 54, 167–183.
- Stahl, C., & Klauer, K. C. (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.
- Vandekerckhove, J., Matzke, D., & Wagenmakers, E.-J. (2015). Model comparison and the principle of parsimony. In J. R. Busemeyer, Z. Wang, J. T. Townsend, & A. Eidels (Eds.), *Oxford handbook of computational and mathematical psychology* (pp. 300-319). Oxford, UK: Oxford University Press.
- Vrieze, S. I. (2012). Model selection and psychological theory: a discussion of the differences between the akaike information criterion (aic) and the bayesian information criterion (bic). *Psychological methods*, 17, 228-243.
- Wagenmakers, E., & Waldorp, L. (2006). Special issue on model selection: Theoretical developments and applications. *Journal of Mathematical Psychology*, 99-214.
- Wickelmaier, F., & Zeileis, A. (2018). Using recursive partitioning to account for parameter heterogeneity in multinomial processing tree models. *Behavior research methods*, 50, 1217-1233.
- Wu, H., Myung, J. I., & Batchelder, W. H. (2010a). Minimum description length model selection of multinomial processing tree models. *Psychonomic bulletin & review*, 17, 275-286.
- Wu, H., Myung, J. I., & Batchelder, W. H. (2010b). On the minimum description length complexity of multinomial processing tree models. *Journal of Mathematical Psychology*, 54, 291-303.
- Yonelinas, A. P., & Jacoby, L. L. (1994). Dissociations of processes in recognition memory: effects of interference and of response speed. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, 48, 516-535.
- Zhang, J. (2012). Model selection with informative normalized maximal likelihood: Data prior and model prior. In E. N. Dzhafarov & L. Perry (Eds.), *Descriptive and normative approaches to human behavior* (p. 303-319). New Jersey: World Scientific.

Appendix

In order to improve performance in terms of the time taken with model fitting, it is advisable to pre-compile the models in C++ and place them in a package that can be loaded in R. This last step is necessary in order to be able to use the pre-compiled functions across multiple cores. Both steps can be achieved conveniently using functions provided by packages `Rcpp` and `devtools` (both require having `Rtools` installed). First, you need to specify a `.cpp` file in which the models are included. The example below, which we will assume to correspond to file `examplemodels.cpp`, includes the SCR model:

```
#include <Rcpp.h>
using namespace Rcpp;

// [[Rcpp::export]]
double SCR(NumericVector Q0, NumericVector chd, NumericVector NN){

  Rcpp::NumericVector Q = pnorm(Q0);
  Rcpp::NumericVector ee(15);

  double Vt = Q[0];
  double Gt = Q[1];
  double Vr = Q[2];
  double Gr = Q[3];
  double a = Q[4];
  double b = Q[5];
  double Vtx = Q[6];
  double Gtx = Q[7];
  double Vrx = Q[8];
  double Grx = Q[9];

  ee[0] = Vt + (1-Vt)*Gt*a + (1-Vt)*(1-Gt)*b*a;
  ee[1] = (1-Vt)*Gt*(1-a) + (1-Vt)*(1-Gt)*b*(1-a);
  ee[2] = (1-Vt)*(1-Gt)*(1-b);
  ee[3] = (1-Vr)*Gr*a + (1-Vr)*(1-Gr)*b*a;
  ee[4] = Vr + (1-Vr)*Gr*(1-a) + (1-Vr)*(1-Gr)*b*(1-a);
  ee[5] = (1-Vr)*(1-Gr)*(1-b);
  ee[6] = Vtx + (1-Vtx)*Gtx*a + (1-Vtx)*(1-Gtx)*b*a;
  ee[7] = (1-Vtx)*Gtx*(1-a) + (1-Vtx)*(1-Gtx)*b*(1-a);
  ee[8] = (1-Vtx)*(1-Gtx)*(1-b);
```

```

ee[9] = (1-Vrx)*Grx*a + (1-Vrx)*(1-Grx)*b*a;
ee[10] = Vrx + (1-Vrx)*Grx*(1-a) + (1-Vrx)*(1-Grx)*b*(1-a);
ee[11] = (1-Vrx)*(1-Grx)*(1-b);
ee[12] = b*a;
ee[13] = b*(1-a);
ee[14] = (1-b);

ee = ee*NN;
double LL = 0.0;
for(int ii=0; ii<15; ii++){
  if(chd[ii] > 0.0){
    LL = LL + chd[ii]*(log(chd[ii])-log(ee[ii]));
  }
}
return LL;
}

```

This model function has three arguments, which should be present in any model function given to function `run_nml`: A parameter vector `Q0`, a frequency-data vector `chd`, and a sample-size vector `NN`. The return value of this function is one-half of the G^2 statistic. Additional models can be included in the same file. Just make sure to add the line `// [[Rcpp::export]]` above each of them.

After specifying the models, they need to be included in a package. This is easily done using function `Rcpp.package.skeleton`, which creates a package folder in your working directory that can then later be installed using function `install`:

```

library(devtools)

Rcpp.package.skeleton("examplemodels",attributes = TRUE, cpp_files = "examplemodels.cpp")

install("examplemodels")

library(examplemodels)

```

Application of NML: Compact Example

```
# load required packages

library(devtools) # required for installing packages from github
install_github("davidkellen/NMLmulti")
library(NMLmulti)


# computes NML penalty for the SCR model

nml_scr <- run_nml(fun=NMLmulti::SCR, parl=10, ks=rep(3,5),
                  Ns=c(rep(100,4),200), fits = 2, cores=6, batchsize=5000,
                  burn=10000, precision=0.1)


# computes NML penalty for the LT model

nml_lt  <- run_nml(fun=NMLmulti::LT, parl=10, ks=rep(3,5),
                  Ns=c(rep(100,4),200), fits = 2, cores=6, batchsize=5000,
                  burn=10000, precision=0.1)
```