

Beyond Curve Fitting? Comment on Liu, Mayer-Kress and Newell.

Andrew Heathcote & Scott Brown

University of Newcastle, Australia

University of California, Irvine

Correspondence should be addressed to:

Andrew Heathcote
School of Behavioural Sciences,
Aviation Building, University of Newcastle
NSW, 2308, Australia
Email: andrew.heathcote@newcastle.edu.au
Phone: 61-2-49215952
FAX: 61-2-216980

Liu, Mayer-Kress and Newell (2003) fit learning curves to movement time data and suggest two new methods for the analysis of learning that they claim go “beyond curve fitting”. Neither their curve fitting nor their new methods take account of measurement noise and so produce inefficient and biased results. We demonstrate these problems using their data, in which variance due to learning is small relative to the level of noise for most participants, and provide better alternatives that are more noise tolerant, more powerful and go “beyond curve fitting” without the extreme bias displayed by Liu et al’s methods.

Liu, Mayer-Kress and Newell (2003; hereafter LMN) examined learning in a discrete movement task, which required participants to produce a five-degree elbow flexion in 125ms. Their data consists of the movement times (MT) for eight participants over 200 trials. Figure 1 shows that changes in MT due to learning were small relative to the level of noise; for participants C, D, G and H only MT for the first trial was greater than later MT s, for E and F only the first two trials were greater, and MT for the first trial was less than on later trials for A. Participant B showed greater learning, but even this was not substantial, with MT s from only the first six trials greater than MT s from later trials. Evidently, any analysis of learning for these data must tolerate substantial noise.

Noisy learning curves, such as LMN's, are usually modelled by an equation of the following form: $MT = f(t) + \varepsilon(\theta)$, where $f(t)$ is a deterministic function of practice trials (t), and $\varepsilon(\theta)$ is a random variable with zero mean (i.e., $E(\varepsilon) = 0$) and parameter vector θ . This form is implicitly assumed by LMN's use of least-squares regression, which is optimally efficient if ε is normal, and asymptotically efficient even when ε is non-normal, under mild regularity conditions (see Jennrich, 1969). However, LMN did not fit MT , but instead fit "absolute error", $AE = |MT - 125|$.

If the aim of fitting is to estimate $f(t)$, least squares regression on AE is clearly inappropriate. When performance is variable, the mean value of AE must be greater than zero, even when mean MT is exactly 125ms. Figure 1 clearly illustrates that performance was variable. The bias induced by using AE is a function of both the level of variability and the distance between $f(t)$ and its asymptote. Larger variability leads to greater average AE values, and as $f(t)$ approaches asymptote the bias increases, systematically distorting the shape of the estimated function. In short, the function estimated by least squares

regression on AE is not $f(t)$ but $f(t) + B(t, \theta)$, where $B(t, \theta)$ is the (trial and noise dependent) bias.

We strongly recommend that researchers avoid the use of AE when their data contain noise, which is always the case for response time measures such as MT (cf. Luce, 1986). In the following sections we report least squares regression results for LMN's data, and analyse the two new methods that they propose, based on MT rather than AE . Our analyses of their two new methods for going "beyond curve fitting" reveal further problems caused by LMN's failure to account for the effects of noise. First, however, we examine LMN's statements about the most powerful method in a researcher's armoury for minimising the effects of noise: averaging.

Averaged Learning Curves

We agree with LMN's caution against averaging learning curves across participants. Brown and Heathcote (2003) provide a simple proof that an arithmetic average of curves has the same form as the component curves if and only if the component curves are linear in all parameters that vary across components. The core of this proof has been known since at least 1821, when Cauchy published it; the interested reader is referred to Aczel (1966) for extensions (e.g., averages other than arithmetic). Heathcote, Brown, and Mewhort (submitted) provide simple graphical and inferential methods of checking whether the linearity condition holds in noisy data. Their analysis of data from Heathcote, Brown and Mewhort's (2000) survey of learning studies suggests that this condition rarely holds in practice. Brown and Heathcote show that averaging across exponential functions that differ in their rate parameters by as little as a single order of magnitude can produce an average curve better fit by a power function. Hence,

curves produced by averaging across participants will usually not be characteristic of participant curves, and any attempt to adjudicate between nonlinear models of participant learning should not rely on data averaged over participants.

However, we disagree with LMNs' caution about averaging over trials. Brown and Heathcote (2003) prove that averaging across trials has *no effect* on the form of exponential functions, producing a change only in the (linear) scale parameter. Because only a linear parameter is changed, the trial average function remains *exactly* exponential and so trial averaging will not cause a bias against the exponential form. Brown and Heathcote also showed that the bias induced by trial averaging of power functions is usually negligible, except in a region of extreme curvature, and even then only when the number of trials averaged is large relative to the extent of that region. Brown and Heathcote's analysis of trial averages of data from Heathcote et al.'s (2000) survey showed that there was little effect on the relative fit of power and exponential functions. Trial averages are very useful because of their ability to reduce noise. For example, Brown and Heathcote's trial averages produced comparable R^2 values to averages across participants with only a minimal risk of the distortion caused by participant averages. Noise causes discrimination between exponential and power functions to fail, approaching no discrimination with extreme noise (Brown & Heathcote, in press). Hence, trial averaging can improve discrimination between curve forms.

Trial averages, and generalizations of trial averages usually referred to as "smooths" (e.g., Bowman & Azzalini, 1997), are particularly useful for exploratory data analysis and graphical display of trends in noisy data, hence their popularity amongst researchers. The thick wavy lines in Figure 1 show broad trial averages (encompassing 50

trials) obtained using the widely available “locally weighted scatterplot smoother” (LOWESS, Cleveland, 1979). The smooth for participant A contradicts LMN’s claim that participant A shows no learning. Although, weak relative to the extent of the noise, there is a clear downward trend in the data, which is confirmed by a statistical test reported in the next section. Participant C, in contrast, shows a clear upward trend, which might arouse suspicion that fatigue effects are present. It is only with the aid of the trial averages that such trends become evident.

The smooths in Figure 1 also illustrate a weakness of trial averages; they are unable to follow the fast changes evident in early trials. Clearly the width of trial averages have to be chosen to suit the rate of change in the data, and trial averages may not be suitable for all data. LMN’s data provides a very difficult case for the analysis of learning both because noise is high and because, for all but participant A, learning is largely confined to the first few trials. However, trial averages remain useful even in this case in order to check for slow variations in the right tail of the learning function that might confound fits of parametric curves, such as the power and exponential, which make a strong assumption of constancy in the right tail.

Curve Fitting

Table 1 provides the proportion of variance accounted for by least squares regressions on the *MT* data for the exponential ($A_E + B_E e^{-rt}$) and power ($A_P + B_P t^{-c}$) functions. These equations have two linear parameters, quantifying the asymptote (A) and the scale of change (B) of the learning function, and one nonlinear parameter, the exponential rate (r) and the power exponent (c). Unlike LMN’s *AE* regressions, the exponential is favoured in the majority of participants. Because of the possible fatigue

effect for participant C revealed by the smooths we used the entire data set (200 trials) rather than LMN's method (using the last 100 trials to estimate A). When only the first 50 and first 10 trials were fit for participant C the fit of the exponential remained superior. For all fits, the proportion of variance accounted for by both curves was significant, with $p < .001$ for all but participant C ($p = .001$ and $p = .047$ for the exponential and power respectively) and participant A ($p = .016$ and $p = .04$ for the exponential and power respectively). The latter result contradicts LMN's statement that participant A showed no learning, and confirms the gradual change made evident by the LOWESS plot in Figure 1. The thin smooth curves in Figure 1 plot the best fitting exponential (solid line) and power (dotted line) functions.

LMN claim that there was "no significant difference in the percentage of variance accounted for" (p. 197) by the power and exponential functions. However, they do not report any inferential tests of this difference, and so have no basis for their statement. Inferential tests can be performed using the nested model technique of Heathcote et al. (2000). This test requires fitting of a four-parameter function that has both power and exponential components, the APEX (Asymptote Power EXponential) function: $A + Be^{-\alpha t} t^{-\beta}$. The APEX function "nests" (i.e., has as a special case) both the power function (when $\alpha = 0$) and exponential function (when $\beta = 0$). Nesting allows the significance of the power and exponential components to be tested using an F test:

$$F_{(f-r, N-f-1)} = \frac{(N-f-1)(R_F^2 - R_R^2)}{(f-r)(1 - R_F^2)}$$

The subscripts F and R refer to the full (i.e., APEX) and reduced (i.e., power or exponential) models with $f = 4$ and $r = 3$ degrees of freedom respectively corresponding to the number of parameters estimated for each model. N is the number of data points.

The results of these tests are presented in Table 1, with $R^2_{APEX}-R^2_E$ testing the power component (i.e., the loss of fit when the power component is omitted) and $R^2_{APEX}-R^2_P$ testing the exponential component (i.e., the loss of fit when the exponential component is omitted). Highly significant evidence for a purely exponential component was obtained for four subjects, whereas the power component was never approached significance, except marginally for participant H. Examination of Figure 1 reveals that, among the participants with highly significant differences, B, E, and F favour the exponential function because of its faster approach to asymptote, and C favours the exponential because the power function overestimates the first point.

Beyond Curve Fitting

We agree with LMN that it is desirable to have tests that go “beyond curve fitting” in the sense of looking at more than relative goodness-of-fit of parametric models. However, we conceptualise this approach, somewhat differently from LMN, as being about non-parametric analyses, that is, analyses that are not prefaced on a particular parametric model or finite set of parametric models. LMN’s “fat tails indicator” (FTI), in contrast, is based on parametric (i.e., power and exponential) model estimates. LMN propose a measure called absolute FTI (λ_{abs}) to discriminate between power and exponential learning curves, and then extend it to a relative FTI (λ_{rel}) test, but “do not claim ... that this ... is optimal in any sense” (p. 203).

To test the bias and efficiency of λ_{abs} , we performed a simulation study using 10000 noisy exponential curves and 10000 noisy power curves, created by adding normally distributed deviates with mean zero and standard deviation 10 to the fits reported by LMN for their participant B (see LMN’s Figure 2)¹. We compared the model

discrimination performance of λ_{abs} against simply picking the model with the best fit (i.e., highest R^2 value). According to LMN's suggestions, we counted $\lambda_{\text{abs}} \leq 0$ as exponential and $\lambda_{\text{abs}} > 0$ as power.

The best-fit test correctly identified power curves in 96% of simulated decisions, whereas λ_{abs} identified the power curves at close to chance levels (48%). For the exponential curves, λ_{abs} was 100% correct and the best-fit test 98% correct. These results indicate a strong exponential bias in λ_{abs} . LMN's *relative* fat-tail indicator (λ_{rel}) is even worse: it is provably biased against power functions. Given a noiseless power function, the power function estimated from the data by any competent method is exactly the same as the data. This equality makes the numerator in LMN's Equation 6 zero, and so λ_{rel} always incorrectly chooses an *exponential* model for pure power curves! We confirmed, using the same simulation methodology applied to λ_{abs} , that this extreme bias also applies to noisy curves; λ_{rel} always classified the power data as exponential. Clearly, LMN's FTI methods should be avoided. In contrast, a simple comparison of goodness-of-fit performs well and we recommend its use.

LMN suggested a second method of going "beyond learning curves" that is truly nonparametric, their discrete proportional error change measure (R_n). R_n is very similar to the "relative learning rate" (RLR) measure which Heathcote, et al. (2000) used to characterise differences and similarities amongst continuous parametric learning curves with apparently unrelated forms. Heathcote et al.'s measure produced simpler results for continuous curves, and the same conclusions apply to R_n , so we use RLR here:

$RLR(t) = -f'(t)/(f(t) - f(\infty))$. The prime indicates differentiation with respect to t , and we assume that f must be once differentiable and strictly monotonic. Constant RLR is a

defining feature of the exponential function, $RLR = r$, so average RLR directly estimates the exponential rate parameter. For a power function RLR decreases to zero hyperbolically ($RLR(t) = c/t$), so detecting a significant decrease in RLR with trials favours a power function, although not uniquely.

Two problems arise when estimating RLR from noisy data. One problem is paramount for asymptotic samples, where the true RLR denominator approaches zero. As the true denominator approaches zero, even very small perturbations due to noise can produce very large fluctuations in RLR estimates. Any attempt to summarise the behaviour of RLR can become dominated by these effects, obscuring useful information about changes in MT with practice in early trials. This behaviour is evident in LMN's Figures 3 and 4, where LMN use a linear regression to summarise the behaviour of the RLR estimates. They report that slopes were not significantly different from zero, and conclude that learning is exponential because RLR does not change with practice. However, it is evident from their Figure 4 that their estimates do not differ significantly from zero overall², because they are swamped by asymptotic noise.

The second problem with LMN's direct calculation of R_n for each pair of trials (and the corresponding calculation of RLR) is that it relies on an accurate estimate of $f(\infty)$. Even a small error in this estimate (as might have happened due to fatigue for participant C) can produce extremely large distortions in RLR estimates. This approach also relies on having a long and stationary set of measurements of asymptotic performance, which is not always available or even practicable. Unfortunately, the two problems tend to compound each other; measurement of asymptotic performance improves the asymptote estimate, but also increases asymptotic noise.

Estimating the *RLR* function via regression rather than direct calculation can solve both problems. For the exponential, $-f'(t) = r(f(t) - f(\infty))$, and so a linear regression with slope r and intercept $-rf(\infty)$ is predicted in a plot of $-dMT(t + \frac{1}{2})$ against $\overline{MT(t + \frac{1}{2})}$, where $dMT(t + \frac{1}{2}) = MT(t+1) - MT(t)$ estimates f' , and $\overline{MT(t + \frac{1}{2})}$ (the average value of MT over trials t and $t+1$), estimates f . Figure 2 shows this plot for participant B. A quadratic regression³ (solid line, Figure 2) revealed a significant linear ($F(1,196)=28.6$, $p<.001$), but not quadratic ($F<1$) component, with $R^2 = 0.129$. It is evident from Figure 2 that only the first five points provide much information about the *RLR* function. The first point is clearly highly influential; when it is removed the linear regression indicated by the dotted line in Figure 2 is obtained, with $R^2 = 0.156$. Solving for the parameters we obtain an asymptote estimate of 124.24 and a rate estimate of 0.3345, both of which are in good agreement with the least squares exponential fit (124.75 and 0.350 respectively).

For the power function, $-f'(t) = c(f(t) - f(\infty))t^{-1}$. Hence, a linear plot of $-dMT(t + \frac{1}{2}) \times (t + \frac{1}{2})$ against $\overline{MT(t + \frac{1}{2})}$ is predicted, with slope c and intercept $-cf(\infty)$. Figure 2 shows this plot for participant H, as they had the strongest parametric evidence for a power function. Neither linear nor quadratic components were significant ($F_s < 1$). The linear slope grossly underestimates the power exponent relative to the least squares power fit (4.824) but only slightly underestimates the asymptote (124.2 vs. 127.0 respectively). These results do not support the power function as an adequate model of learning, and indicate that one slow point (calculated from trials 1 and 2) and two fast points (calculated from trials 3 and 4 and 4 and 5) are highly influential.

Finally, if the aim is simply to investigate possible deviations from a reference curve form (e.g., exponential), averaging across participant *RLR* function estimation plots (e.g., over $-\Delta MT(t + \frac{1}{2})$ and $\overline{MT(t + \frac{1}{2})}$ for the exponential) is both convenient and mathematically appropriate, due to their linearity. Averaging will distort the exact form of the non-linearity in the plot for functions of *different* forms, but deviation from linearity will still be evident and the reference curve can be rejected without confounding. Note, however, that this approach assumes each participant has the same curve form, differing only in parameters.

Given that earlier tests favour the exponential for most subjects, we calculated an average *RLR* function estimation plot for the exponential (Figure 2). A quadratic regression gave $R^2 = 0.206$, with a significant linear ($F(1,196)=47.8, p<.001$) but not quadratic ($F(1,196)=2.57, p=.11$) component. The best fitting linear regression line ($R^2=0.195$) has a slope that only slightly overestimates the geometric mean of the rate parameters from the least-squares exponential fits (0.412) and the arithmetic mean of their intercepts (127.7 and 126.8 respectively). Given the weakness of learning and high noise levels for most participants, the agreement is surprisingly good. However, it must be acknowledged that nonparametric approaches pay a cost in power relative to parametric approaches, at least when the assumptions of the later are true. In LMN's data, which has both little information about the systematic structure of the learning curves and high noise levels, only participant B provided clear results with *RLR* plots, and even in this case an influential point had to be censored.

Brown and Heathcote (2002) suggested an alternative nonparametric approach⁴, which is more powerful because it makes a slightly more restrictive, but still plausible,

assumption about the form of the learning curve: that it is smooth. This method uses trial averaging techniques (smooths) and bootstrap estimation (e.g., Davison & Hinkley, 1997) to attach probability values to hypotheses of the form “this regression curve is not significantly different from the data generating model”. It extends Azzalini, Bowman and Hardle’s (1989) approach to model selection by compensating for smoothing bias. Confidence intervals can be obtained in a completely non-parametric manner (i.e., without assuming any parametric form for the noise, cf. Hardle & Marron, 1991) using a Wild bootstrap.

We applied Brown and Heathcote’s (2002) method to LMN’s data. These analyses returned the inferential probability values shown in Table 2. Higher probability values indicate greater evidence in favour of the corresponding model. Comparison of power and exponential model probabilities are generally consistent with goodness of fit, although this is not necessarily so (see Brown & Heathcote for an example where they clearly disagree). Figure 3 shows the smooths and corresponding confidence intervals for participants B and H. Table 2 indicates that LMN’s data does not have sufficient power to reject either model, except for these two participants. Both models perform fairly poorly for participant H, and we conclude that neither is entirely adequate in this case. For participant B, Figure 3 shows that the power function badly underestimates early performance, and overestimates later performance. The exponential model shows some overestimation around trial 75, but is generally much better. Figure 3 illustrates the value of confidence intervals generated from Brown and Heathcote’s technique – the actual *type* of misfit can be identified, even in noisy data.

General Discussion

LMN suggested that traditional methods of discriminating between learning curve models are inadequate, and proposed two new methods to go “beyond curve fitting”. However, their implementation of curve fitting, using absolute error rather than the raw movement time data, produces biased results. When we corrected this error we found that, despite the high noise levels and rapid learning in LMN’s data, properly constructed and efficient statistical tests do clearly adjudicate in favour of an exponential function for half of LMN’s participants. For the remaining participants these tests also convey important information: that the data are too noisy to provide clear evidence either way.

LMN’s mistake in assuming that participant A shows no significant learning illustrates both the importance of both formal statistical testing and of trial averages in detecting slow trends obscured by high levels of noise. We also reviewed evidence that appropriate trial averages cause little or no distortion in curve form for the sorts of smooth models considered by LMN. Hence, although we agree that participant averages should be avoided, we differ from LMN in recommending trial averaging as a useful method for investigating learning curves.

Neither of LMN’s two methods for going “beyond curve fitting” are useful. The fat-tails indicator (FTI) is extremely biased in favour of an exponential model. Therefore, it is no surprise that this method provided “clearer” evidence in favour of the exponential model than curve fitting. Their R_n measure does not suffer from bias, but is extremely intolerant of noise and has little power to discriminate between curve forms. LMN claimed that failure to find significant decreases in R_n with practice favour an exponential

function. This conclusion is invalid because it affirms the null hypothesis with a test that demonstrably has no power.

Our analyses have brought to light serious flaws in LMN's methods. However, better methods are readily available. In the domain of curve fitting, the standard approach of choosing the model with best least-squares fit vastly outperformed the FTI method, and Heathcote et al.'s (2000) nested model testing method was shown to provide a powerful inferential test of curve form. Brown and Heathcote's (2002) method, and the linear *RLR* plot method suggested here, fulfil LMN's desire to go "beyond curve fitting" in that they provide information about distinctive properties of, and systematic deviations from, learning curve models in a truly nonparametric manner. Despite the fact that our conclusions about their movement time curves do not differ from LMN (we support an exponential model for most participants) we strongly recommend that researchers do not adopt their methods.

References

- Azzalini, A., Bowman, A.W., & Hardle, W (1989) On the use of non-parametric regression for model checking. Biometrika, 76, 1-11
- Aczel, J. (1966) Lectures on functional equations and their applications. London: Academic Press.
- Bowman, A.W. & Azzalini, A. (1997) Applied smoothing techniques for data analysis: The kernel approach with S-Plus illustrations. Oxford: Clarendon Press.
- Brown, S., & Heathcote, A. (2002). On the use of Nonparametric Regression in Assessing Parametric Regression Model. Journal of Mathematical Psychology, 46, 716-730.
- Brown, S. & Heathcote, A. (2003). Averaging learning curves across and within participants. Behaviour Research Methods, Instruments & Computers, 35, 11-21.
- Brown, S. & Heathcote, A. (in press). Bias in Exponential and Power Function Fits Due to Noise: Comment on Myung, Kim and Pitt, Memory and Cognition.
- Cleveland, W. S. (1979). Robust locally weighted regression and smoothing scatterplots. Journal of the American Statistical Association, 74, 829–836.
- Davison, A.C. & Hinkley, D.V. (1997) Bootstrap methods and their application. Cambridge: Cambridge University Press.
- Heathcote, A., Brown, S. & Mewhort, D.J.K. (2000) Repealing the power law: The case for an exponential law of practice. Psychonomic Bulletin and Review, 7, 185-207.
- Heathcote, A., Brown, S. & Mewhort, D.J.K. (submitted). Average Curves and Repeated-Measures Analysis of Variance.

Hardle, W. & Marron, J.S. (1991) Bootstrap simultaneous error bars for non-parametric regression. Annals of Statistics, 19, 778-796

Jennrich, R. I. (1969). Asymptotic properties of nonlinear least squares estimators. Annals of Mathematical Statistics, 40, 633-643.

Luce, R. D. (1986). Response Times. New York: Oxford University Press.

Liu, Y-T, Mayer-Kress, G. & Newell, K. M. (2003). Beyond curve fitting: A dynamical systems account of exponential learning in a discrete timing task. Journal of Motor Behavior, 35, 197-207.

Wand, M. P., & Jones, M. C. (1995). Kernel Smoothing. London: Chapman & Hall.

Footnotes

¹ This participant was chosen because they displayed the clearest evidence of learning in Figure 1. Apart from an overall decrease in discrimination, a similar pattern of results was obtained using a wide range of learning function parameters and levels of noise (SD = 5, 10, 15, 50), and similarly for LMN's estimated parameters from Participants A and D.

² By definition any decreasing function must have $RLR > 0$, although this might be difficult to detect in a power function as RLR approaches zero with practice. Hence, failure to find RLR estimates significantly greater than zero could be interpreted as evidence against an exponential learning curve, as its RLR remains greater than zero even with extended practice. We prefer an interpretation in terms of lack of power because of the very large fluctuations evident in RLR estimates.

³ Higher order polynomials, or alternative sets of basis functions, can also be used. Strictly, univariate regression is inappropriate, as $\overline{MT}(t + \frac{1}{2})$ is measured with error. However, $\overline{MT}(t + \frac{1}{2})$ necessarily has less error than $\Delta MT(t + \frac{1}{2})$, and so this approach is approximately correct and the approximation is convenient as univariate regression software is widely available.

⁴ Matlab code to perform this analysis is available from http://www.newcastle.edu.au/school/behav-sci/ncl/software_repos.html.

Tables

Table 1. Fits of the exponential (subscript E), power (subscript P) and APEX (subscript APEX) functions to *MT* data from all 200 trials for each participant (A..H). The F ratios, $df=(1,195)$, test the R^2 difference in the previous column.

	R^2_E	R^2_P	R^2_{APEX}	$R^2_{APEX}-R^2_E$	F	p	$R^2_{APEX}-R^2_P$	F	p
A	0.051	0.041	0.051	0.000	0.00	0.997	0.010	2.01	0.157
B	0.843	0.733	0.843	0.000	0.00	0.996	0.110	136.05	0.000
C	0.082	0.040	0.082	0.000	0.00	1.000	0.042	9.00	0.003
D	0.098	0.097	0.098	0.000	0.00	1.000	0.001	0.23	0.633
E	0.398	0.356	0.398	0.000	0.00	1.000	0.042	13.61	0.000
F	0.357	0.309	0.357	0.000	0.00	1.000	0.048	14.49	0.000
G	0.167	0.170	0.170	0.003	0.63	0.428	0.000	0.00	1.000
H	0.346	0.359	0.358	0.011	3.46	0.064	0.000	0.00	1.000

Table 2. Inferential probability values from Brown and Heathcote's (2002) method.

	A	B	C	D	E	F	G	H
Exponential	.504	.044	.623	.677	.594	.073	.242	.038
Power	.452	.003	.622	.692	.288	.070	.275	.034

Figure Captions

Figure 1. LMN's *MT* (movement time) data (dots) for each participant (A...H, see ordinate labels), LOWESS smooths (dash-dot lines) and best fitting exponential (solid lines) and power (dashed lines) functions. The target *MT* was 125 ms (horizontal dotted lines).

Figure 2. Exponential (participant A, participant average) and power (participant H) *RLR* function estimation plots. Dashed lines indicate a best fitting linear regression and solid lines indicate a best fitting quadratic regression.

Figure 3. Nonparametric regression estimates for Subjects B and H (solid lines), along with pointwise 95% confidence intervals (dotted lines) on the location of these estimates under the assumption of either power or exponential models. Error structure was modelled using 5000 bootstrap samples from the residuals; smooths were calculated using local linear regression with an Epanechnikov kernel of half-width $h=2$.





Figure 1 (continues next page)





Figure 1 (continued)





Figure 1 (continued)





Figure 1 (continued)


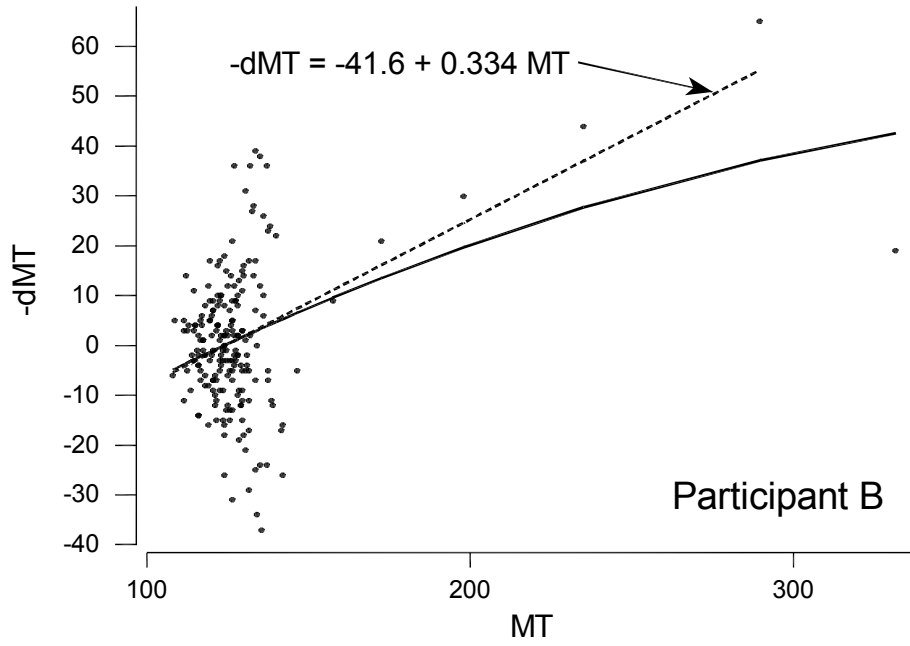


Figure 2 (continues next page)

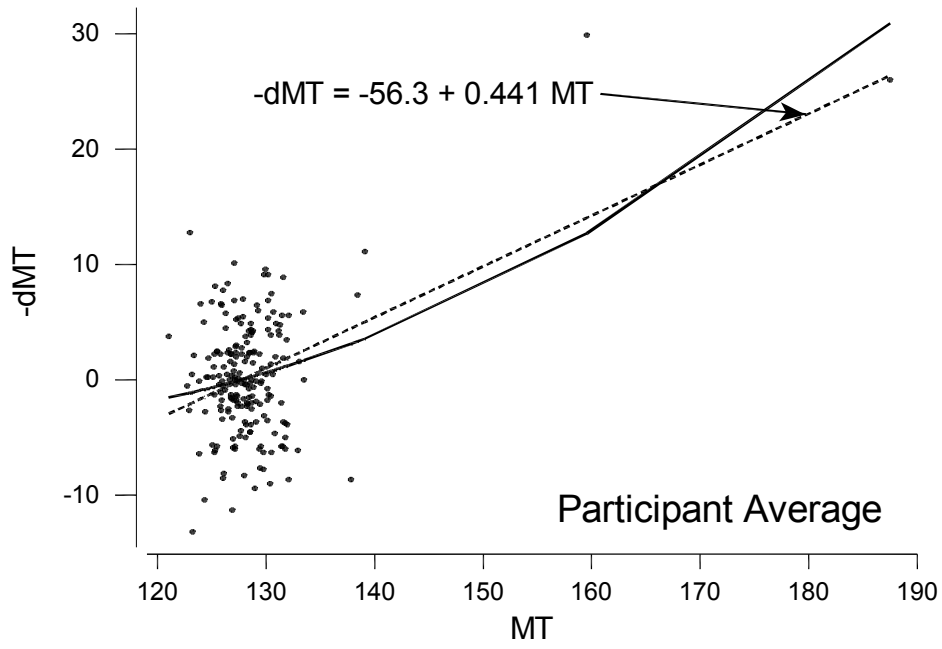


Figure 2 (continued)


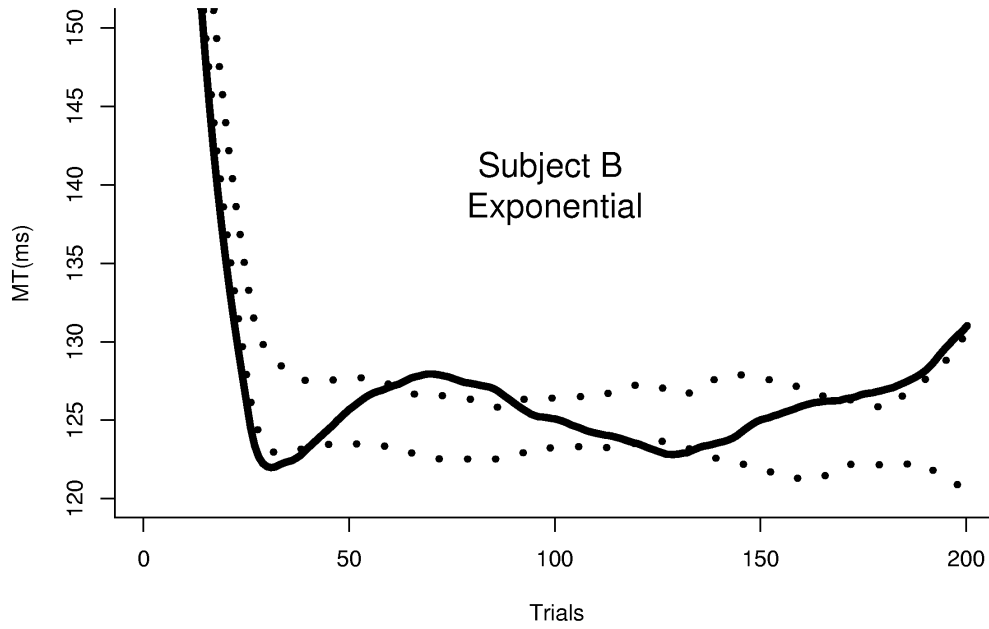


Figure 3 (continues next page)


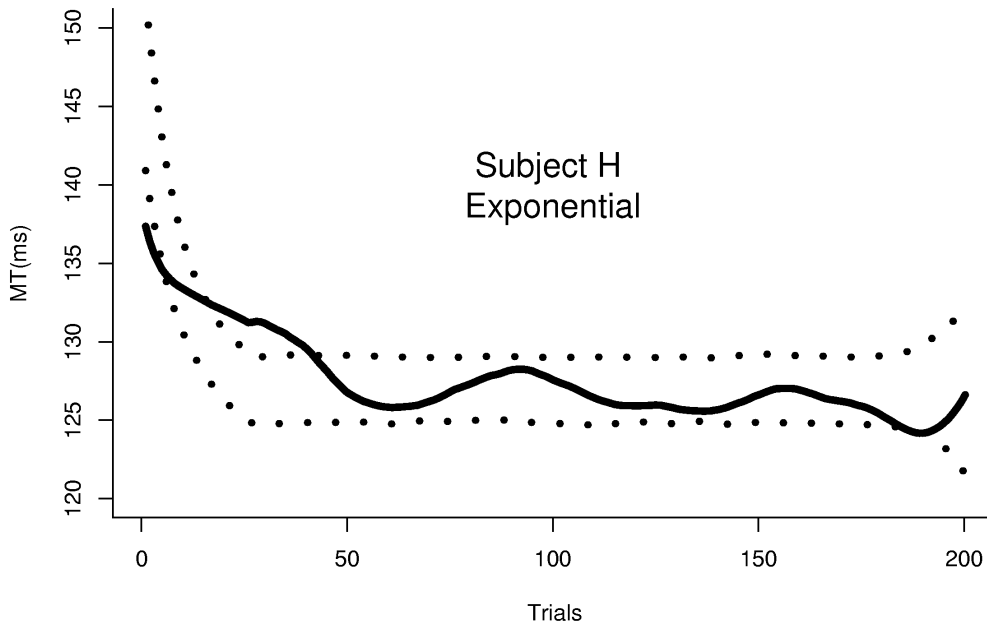


Figure 3 (continued)