DOCUMENT CONTROL SHEET

	ORIGINATOR'S REF.		SECURITY CLASS.
	NLR-TP-2003-339		Unclassified

ORGINATOR

National Aerospace Laboratory NLR, Amsterdam, The Netherlands

TITLE

The acquisition of complex skills and the linear rate model

PRESENTED AT

The colloquium of the Stochastics and Optimization Group, Department of Mathematics, Utrecht University, 15 January, 2003.

AUTHORS	DATE	PP	REF
J.J.M. Roessingh, J.J. Koenderink and A.M.L. Kappers	July 2003	44	40

DESCRIPTORS

COMPLEX SKILL ACQUISITION
LEARNING
SELECTION
LINEAR RATE MODEL
LEARNING CURVE
POWER LAW OF PRACTICE,
TRAINING
SPACE FORTRESS GAME
PARAMETER ESTIMATION

ABSTRACT

Since the beginning of the 20th century, smooth functions have been proposed to describe or predict the learning curve of individuals as a function of practice time or practice trials. In contrast, we propose a 'progressive average function' to describe the learning curve for complex speed-based skills. In the absence of random trial-to-trial variability, a progressive average function with one free parameter corresponds exactly to the ill-specified power law of practice, proposed by de Jong (1957) and Newell and Rosenbloom (1981).

We measured the individual skill level of trainees on the basis of times needed to successfully complete subsequent trials on a complex computer game. Each trainee received 16 hours of practice training. In all we measured 13,000 trial times. We compared the progressive average function with regular learning curve models with respect to their ability to describe experimental learning data. The progressive average function turned out to provide the most parsimonious and consistent description of the data. According to the linear rate model, on which this function is based, the times taken to complete trials on a complex task are stochastic variables that come from a long-tailed (approximately log-Gaussian) distribution. In the long-term, these stochastic variables are controlled by a deterministic trend, which is the result of a naïve learning strategy. The model predicts stability of performance with practice, long-lasting individual differences and specificity of skill-transfer.

Nationaal Lucht- en Ruimtevaartlaboratorium
National Aerospace Laboratory NLR

NLR-TP-2003-339

The acquisition of complex skills and the linear rate model

J.J.M. Roessingh*, J.J. Koenderink** and A.M.L. Kappers**

*National Aerospace Laboratory NLR, **Helmholtz Institute, Utrecht

This report is based on a colloquium for the Stochastics and Optimization Group of the department of Mathematics at Utrecht University, The Netherlands on 15 January 2003.

Customer: National Aerospace Laboratory NLR
Working Plan number: V.1.C.2
Owner: National Aerospace Laboratory NLR
Division: Flight
Distribution: Unlimited
Classification title: Unclassified
July 2003

Approved by author: 23/07	Approved by project manager: 23/7	Approved by project managing department: 23/7

Summary

Since the beginning of the 20th century, smooth functions have been proposed to describe or predict the learning curve of individuals as a function of practice time or practice trials. In contrast, we propose a ‘progressive average function’ to describe the learning curve for complex speed-based skills. In the absence of random trial-to-trial variability, a progressive average function with one free parameter corresponds exactly to the ill-specified power law of practice, proposed by de Jong (1957) and Newell and Rosenbloom (1981).

We measured the individual skill level of trainees on the basis of times needed to successfully complete subsequent trials on a complex computer game. Each trainee received 16 hours of practice training. In all we measured 13,000 trial times. We compared the progressive average function with regular learning curve models with respect to their ability to describe experimental learning data. The progressive average function turned out to provide the most parsimonious and consistent description of the data. According to the linear rate model, on which this function is based, the times taken to complete trials on a complex task are stochastic variables that come from a long-tailed (approximately log-Gaussian) distribution. In the long-term, these stochastic variables are controlled by a deterministic trend, which is the result of a naïve learning strategy. The model predicts stability of performance with practice, long-lasting individual differences and specificity of skill-transfer.

This page is intentionally left blank

Contents

(43 pages in total)

1 Introduction

1.1 Approach and goal

Throughout this study we use the term 'trial time' rather than 'response time', to denote the time needed to successfully complete a trial on a task. The term 'response time' is commonly associated with more basic laboratory tasks, while this research deals with a complex task, which requires a series of motor responses to complete a trial on the task.

Crossman (1959) noted that the overall frequency histogram of a series of trial times becomes increasingly skewed with practice and that the variability of trial times decreases with practice. To measure this process in detail, we investigated a representative speed-task, a thoroughly investigated computer game. This speed-task allows skill level to be measured principally by the trainees' speed of performance, as soon as success (completion of the task) can be taken for granted, and thus prevents trade-off effects between speed and accuracy in task performance. In order to make a detailed investigation of the distribution of trial times in a learning process we collected large numbers of trial times (on average 2175 trial times for each of the six participants) over a relatively long training period (16 hours total practice time, i.e. 'cumulative response time'). We also noted the precise point at which these trial times were measured in the individual learning process. To allow for statistical time series analysis, we ensured that the sum of trial times exactly equaled total practice time and thus that the increase in skill level gained through practice was fully captured by the measured trial times.

With these data we attempt to identify the change process in the distribution of trial times. This change process can be modeled by a function with only one free parameter that will predict the learning curve of individuals when acquiring complex skills.

1.2 Complex skills

Schneider (1985) typified complex skills, commonly called 'high performance skills', in the following terms: (1) the skills can only be acquired after many hours of practice, sometimes even hundreds of hours; (2) the skills cannot be mastered by a substantial number of potential trainees; (3) the skills of expert operators are qualitatively superior to those of novices.

We describe some characteristics of the complex skills that pilots need to handle a serious equipment failure with an aircraft. This is a typical training scenario in a flight simulator. In such a scenario the core task would invariably be to stabilize the aircraft in terms of attitude and speed and to maintain a safe altitude; such a task requires both expertise and flying skills. Task management would require prioritization, taking into account a range of possible subtasks, for example, selecting an appropriate place to land, crew coordination, jettisoning weight, emergency radio handling, configuring the aircraft for landing, etc.

In general, complex skills are needed to perform a complex task, namely, a set of diverse core tasks, such as the continuous task of controlling a vehicle, support tasks, and managerial tasks. Those diverse tasks, in their turn, can be decomposed into subtasks or part-tasks. Some of these part-tasks will have to be executed simultaneously, requiring so-called time-sharing ability on the part of the human operator. Using a subset of possible methods, the execution of a trial on the whole task begins with low-level part-task execution and moves on towards higher-level part-task-execution. The human operator requires tactics to counter short-term threats, and strategies that take into account longer term effects, such as expended resources and overall safety of the operation. Some of these methods are clearly less effective than other; Learning to select the most effective methods of attaining the top-level goal of the whole task is at the heart of the learning process.

1.3 Choice of a complex task

In order to collect the skill acquisition data, we used the Space Fortress (SF) Game as our experimental task. This is a Personal Computer game that was specifically designed for the study of complex skills learning and essentially follows the lines that we sketched in the previous section. Skill acquisition with the SF game was extensively examined in experimental research. This research has been documented in a special volume of Acta Psychologica (edited by Donchin, Fabiani & Sanders, 1989) and in various later research publications. The statement that SF is a representative complex skill trainer is supported by field studies at flight schools where SF has been used in flight training (e.g. Gopher, Weil & Bareket, 1992, 1994, Hart & Battiste, 1992, Vidulich, McCoy & Crabtree, 1995).

In our study the instruction to the trainee is basically to “destroy the space fortress as quickly as possible”. This means that the trainee has to control the space ship, which has intricate control behavior, and has to bring it to the appropriate position for attacking the space fortress. The task further requires detection, recognition and identification of threats (e.g. moving mines), which have to be neutralized. The game also imposes cognitive requirements since various procedures have to be adhered to and scarce resources such as ammunition have to be managed. The game clearly calls for tactical and strategic skills. For example, participants in a large-scale control group experiment organized by Foss, Fabiani, Mane & Donchin (1989) could be ordered on the basis of the seven different strategies that they used to control the space ship. The SF game provides good research material in that it is sufficiently complex and interesting to guarantee a very long skill acquisition process and at the same time allows for reliable measurement of valid performance measures, in our case thousands of subsequent trial times per trainee.

1.4 Overview of the study

Until recently the ‘power law of practice’ was considered as a fundamental benchmark result in psychology (Palmeri, 1999) and was included in the handbooks of the applied scientist (e.g. Boff & Lincoln, 1988, pp. 860-863). Experimenters fitted the power law of practice to complex skill data for diagnostic applications, at the assumption that the lawfulness implied by the word ‘law’ made it a standard diagnostic tool, generally applicable in situations for which it was formulated (foremost by Newell & Rosenbloom, 1981). However, problems with the power law frustrate meaningful application (see Heathcote, Brown & Mewhort, 2000). Therefore, we address three important issues connected to the power law of practice in the next sections.

We first consider the ‘overfitting’ problem, that is, the problem that a power law with three or four free parameters can fit a series of trial times from an individual trainee in many different ways. To constrain the parameters in a meaningful way, we use an asymptotic power law result (Györgyi & Tishby, 1990) with respect to learning of neural networks.

Second, we demonstrate that power laws are disguised 'progressive average functions', which, we think, clarifies the 'power law artifact' (Myung, Kim & Pitt, 2000). Moreover, progressive average functions are a more parsimonious description of learning data.

Third, we consider the distribution of random trial-to-trial error, which has implications for the technique with which learning curves for complex skills should be fitted to series of trial times.

Thereafter, we describe the Space Fortress experiment. We start the analysis of results with an examination of the validity of trial times relative to other performance scores for the task at hand. We then analyze the resultant empirical learning curves of the six trainees that participated in the experiment. We subsequently compare competing models (notably, power functions, exponential functions and the progressive average function, the latter function derived from the 'linear rate model'). Finally we investigate the distribution of trial times and provide a discussion of results.

2 Theory

2.1 The shape of practice curves

The two most conspicuous features of long series of trials on a speed-based task are: (1) the decreasing trend in the series of trial times, and (2) the slowing down of this decreasing trend with practice. The learning curve literature (e.g. Thurstone, 1919, Restle & Greeno, 1970, Mazur & Hastie, 1978) provides several quantitative models that have been used in attempts to explain or describe the supposed reduction in trial times, either as a function of practice time or as a function of practice trials.

Mazur & Hastie (1978) reviewed practice data of 23 experiments that dealt with a wide range of different skills. They examined the shape of the learning curve by contrasting two mathematical functions: (1) a function that modeled the learning process as a replacement process, in which incorrect response tendencies are replaced by correct ones, and (2) a function that modeled the learning process as a process of accumulation, in which incorrect response tendencies remain constant and correct response tendencies increase with practice. Their major conclusion was that the accumulation model describes the shape of most learning curves better than does the replacement model.

2.2 Power laws

To describe the expected reduction in trial time $\boldsymbol{T}_n$ as a function of the number of completed trials n, a power function has been proposed by many researchers:

$$E[\boldsymbol{T}_n] = b\,n^{-\alpha}\,. \qquad (1)$$

$E[\boldsymbol{T}_n]$ is the expected value of trial time $\boldsymbol{T}$ on trial n. The boldface notation indicates that $\boldsymbol{T}$ is a random variable. When referring to observed trial times (i.e., to samples from $\boldsymbol{T}$), we will use the notation T. The power law has two free parameters, a scale parameter b and a shape parameter (exponent) α. It includes a hyperbolic function as a special case (when the exponent α=1).

Snoddy (1926) observed the power law regularity when he plotted the trial times against trial number for practice trials of subjects that performed a perceptual-motor task (mirror-tracing) in a clinical study. Snoddy used double-logarithmic paper and found that the data points were closely following a straight line.

Later, de Jong (1957) provided a review of some data sets on industrial tasks, using Snoddy's law with an extra parameter to account for the time consumed by the incompressible part of the task (when the human operator had to wait for the machine during a trial). Following de Jong, Crossman (1959) proposed a model that described various data sets, such as the data from a study of cigar production by machine operators. Crossman explained the observed reduction in trial times in terms of the acquisition of more efficient methods to do the task, a method being defined as a "distinguishable action pattern". The student operator was assumed to possess a fixed repertoire of such methods, from which he picks one by chance at each trial. The use of such methods would then affect the "habit strength", defined as the probability of use of a method. Crossman's assumption that the student operator possesses a fixed repertoire of methods, rather than a growing repertoire of methods, makes his model a replacement model, rather than an accumulation model. However, Crossman's model did not explain 'de Jong's law'. Although Crossman could not obtain a closed-form functional description for the learning curve predicted by his model, his graphs revealed that it was clearly not a power law.

Newell & Rosenbloom (1981) provided a review of eleven data-sets, that addressed a wide range of tasks and yielded superior fits for a power law over other investigated functions, such as an exponential function of trial number *n*. They proposed a four-parameter power law of practice:

$$E[\boldsymbol{T}_n] = T_d + b \cdot (n + n_0)^{-\alpha}, \qquad (2)$$

with two additional free parameters, dead time T_d, and shift parameter n_0. The first extra parameter T_d was to account for an asymptotic trial time, when trial number goes to infinity. According to Newell & Rosenbloom (1981), n_0 can be interpreted as number of trials of learning that occurred prior to the first trial as measured and thus 'identifies the true starting point of learning'. Also, the parameter *b* can be interpreted as the performance time of the first trial (when n=1 and n_0=T_d=0). The parameter α can be interpreted as the learning rate and is thought to take a value that is less than 1.

2.3 Overfitting

Since Newell & Rosenbloom's influential chapter on the power law, it has been a commonly investigated model for the learning curve in complex tasks (e.g. recently, Lee & Anderson, 2001, Speelman & Kirsner, 2001). Today researchers are once again discussing whether a power function or an exponential function is the best fitting function to skill acquisition data. It has been argued that the power law proposed by Newell & Rosenbloom is an artifact of averaging (Anderson & Tweney, 1997, Myung, Kim & Pitt, 2000). When power laws were compared with other smooth functions for the learning curve, averaging of trial times over trainees or trials works in the advantage of a power law. Without averaging over trials or trainees, an exponential function often seems to fit better than a power function.

Heathcote, Brown & Mewhort (2000) extensively documented the case of a power function against an exponential. On the basis of re-analysis (with improved curve fitting and data handling techniques) of experimental data sets, including a review of some of the data sets originally analyzed by Newell & Rosenbloom, they conclude that an exponential function provides a better fit than a power law for the majority of data sets.

However, in addition to the supposed 'power law artifact' as a result of averaging (which will be explained in the next section), it should not be overlooked that the power law of equation 2 can be rewritten as:

$$E[\boldsymbol{T}_n] = T_d + b \cdot n_0^{-\alpha} \cdot n^{-\alpha'},\ with\ \alpha' = \alpha \cdot \frac{\log[\frac{n+n_0}{n_0}]}{\log n}. \qquad (3)$$

This expression is algebraically equivalent to equation 2, which can be verified by moving T_d to the left hand side and taking the logarithm at both sides. Equation 3 shows, that except for a term T_d, the expected trial time E[$\boldsymbol{T}_n$] is contingent on three multiplicative factors, in the same sense that the volume of a box is contingent on its length, width and depth. To verify the law for the volume of a box, we need to fix two of the variables, for example, width and depth, while determining the relationship between length and volume. Likewise, to verify equation 3 (or equation 2), we should fix the value of two factors, for example the factors b and $n_0^{-\alpha}$, and determine the relationship between and $n^{-\alpha`}$ and $\boldsymbol{T}_n$.

Thus, equation 3 demonstrates that variations in the unknown parameters b, n_0, α and trial number n are not separable. In other words, the parameters share redundant information about the trial time $\boldsymbol{T}_n$. These parameters will usually be estimated from a series of n empirical trial times of an individual, yielding a time series $\{T_1, .. , T_n\}$ with random error. Since many combinations of parameter values fit the data equally well or bad, an arbitrary balance between parameter values will be chosen, depending on the nature of random error and the estimation method employed. The estimation method will thus produce non-interpretable results. As a result of an attempt of fitting equation 2 to a large number of data sets, Heathcote et al. noted: '.. it is likely that parameter estimates from the power function are unrelated to the psychological processes underlying learning or, at best, related in a complex way determined by the best imitation that a fitting algorithm can find.' This is indeed true and can be illustrated with an example. In an analysis of learning a complex discrimination task, Seibel (1963, p. 220) fitted the power law of equation 2 to a series of reaction times for subject JK. Seibel estimated the following parameters: T_d=0.3, b=3 s, n_0=2, α=0.9.

However, when Newell & Rosenbloom (1981, p. 25) fitted the same law to the same data for subject JK, using a different estimation method, they estimated: T_d=0.3, b=2440 s, n_0=2690, α=0.95. Clearly, different methods of curve fitting caused a change in the parameter estimates for b and n_0 by several orders of magnitude. Moreover, when Newell & Rosenbloom (p. 4) fitted the simple power law of equation 1 to the same data of Seibel's subject JK, they estimated b=12 s and α=0.32. In all three cases an R^2 (percentage of variance accounted for) of more than 99% was reported. While Seibel (1963) noticed that his fits were unreliable, Newell & Rosenbloom (1981) largely ignored the problem in their re-analysis.

In summary, some prior knowledge on the values of two of the parameters b, n_0, α is a necessity, if one wants a sensible estimate the third parameter. For example, when α and n_0 are assigned constant values, that is, on the basis of theoretical assumptions, can we find a sensible estimate for b. If one wants to predict the trial times of a trainee that is 'completely' new to a task, the 'prior practice' parameter n_0 does not seem to pose a large problem, and can be set to zero or an appropriately low value, but this does not resolve the inseparability of b and α.

2.4 Constraining the power law in the exponent α

Although there are many psychological theories that predict a power law as a function of trials (e.g. Newell & Rosenbloom, 1981, Anderson, 1984, Logan, 1988a, 1992, Nosofsky & Palmeri, 1997), there are few psychological theories that make a theoretical prediction on the value of α. In fact, every conceivable value between zero and one has been found in empirical data of human operators, irrespective of the type of task, which may be largely attributable to the problem of inseparable parameters.

However, in recent years, the problem of learning from examples has been an attractive topic in statistical mechanics (see Engel & van den Broeck, 2001). In this context, the learning process corresponds to setting the coupling parameters between neural units in a feed-forward network, such as a perceptron (Rosenblatt, 1961) with a large number of neural units, such that the network stores the desired pattern. The network has to implement a functional relation, i.e. a target rule, defined on all possible input patterns. When presented with a novel example (a noisy pattern) it should provide a correct prediction of the desired pattern. If those patterns are known only through some noisy examples, then the network should 'inductively infer' the correct patterns, that is, generalize (Györgyi & Tishby, 1990). Thus, the aim of learning is generalization. The learning curves of generalization error, which is the probability of false prediction on a novel example, have been calculated for neural networks of different complexity. Universal behavior in the learning curve of neural networks has been calculated for the case that the number of examples gets large, irrespective of the specific architecture of the network. The learning curve is a power function of the number of presented examples. Györgyi & Tishby (1990) calculated that the value of the exponent α depends on the complexity of the task:

$$\alpha = \frac{1+\delta}{1+3\delta}, \qquad (4)$$

in which δ is a measure for stochastic noise at the input of the network and can have a value between zero and one. For the simplest tasks (δ=0), the target rule can be easily extracted from the examples and this rule seems deterministic and realizable by the learner. Thus, for the case δ=0, the exponent α=1, and learning follows a hyperbolic function. For the most difficult tasks (δ=1), when the target rule seems unrealizable or stochastic to the learner, the exponent α=½.

However, there is debate on the exact value of the exponent α. By using an analogy between the fluctuation of the training error and the motion of a random walk, when training a network with noisy examples, Kabashima & Shinomoto (cited in Uezu & Kabashima, 1996) calculated a different result for the power law exponent of the learning curve of generalization error, namely:

$$\alpha = \frac{1+\delta}{1+2\delta}. \qquad (5)$$

This expression gives a power law with exponent $\alpha={}^2/_3$ for the case that the target rule seems unrealizable and stochastic to the learner (i.e. when δ=1). Uezu & Kabashima re-calculated the learning curve for this particular case using a statistical mechanical method (but different to that of Györgyi & Tishby) and found a solution for α that is consistent with Kabashima & Shinomoto's result ($\alpha={}^2/_3$) up to a logarithmic correction.

2.5 A linear rate model for learning

Much of the learning theory, old and new, has been expressed in terms of rate models (e.g. Restle & Greeno, 1970, Mazur & Hastie, 1978, Gallistel & Gibbon, 2000), in the sense that rates are the reciprocals of the random time intervals between learning events. In this section we use a linear increasing rate to model learning as accumulation.

In the current research we use a speed-based criterion, which means that the accuracy (or likewise the error in the result) with which the task must be accomplished is set. Performance can thus be measured by the rate with which subsequent trials on the task are completed successfully. We propose a simple quantitative model for complex skill acquisition that assumes that a trainee's performance (i.e. the rate with which subsequent trials are completed) increases linearly with practice time, and we show that this yields a power law with exponent $\alpha=\frac{1}{2}$.

We generalize the model to cater for dead time, which is inevitable in any task requiring motor responses. In building a model for the learning curve, we will use the random variable $\boldsymbol{t}$ for practice time (cumulative trial time). For practice time after n trials we will use $\boldsymbol{t}_n$, which is the sum of a series of n trial times $\boldsymbol{T}_i$ (i=1..n). Hence, by definition, the n^{th} trial time $\boldsymbol{T}_n$ is the interval between $\boldsymbol{t}_n$ and $\boldsymbol{t}_{n-1}$:

$$\boldsymbol{T}_n = \boldsymbol{t}_n - \boldsymbol{t}_{n-1}\,. \tag{6}$$

We assume that each trial time contains an incompressible dead time T_d in which practice stands still. Total practice time $\boldsymbol{t}_n$ consists of n trials, and thus includes a timeframe with duration $n \cdot T_d$ with no practice and hence no increase in performance. We define an auxiliary random variable for 'corrected practice time':

$$\boldsymbol{s}_n = \boldsymbol{t}_n - n \cdot T_d\,. \tag{7}$$

We also define a random variable $\boldsymbol{\lambda}_n$ for 'corrected rate' which is the number of trials Δn that is completed in a time interval $\Delta \boldsymbol{s}_n$. Thus, according to these definitions, corrected rate at the n^{th} trial is defined as:

$$\boldsymbol{\lambda}_n = \frac{\Delta n}{\Delta \boldsymbol{s}_n} = \frac{1}{\boldsymbol{s}_n - \boldsymbol{s}_{n-1}} = \frac{1}{\boldsymbol{T}_n - T_d}\,. \tag{8}$$

When we assume corrected rate to increase linearly with corrected practice time, we obtain a model for the expected rate at which subsequent trials are completed:

$$E[\boldsymbol{\lambda}_n] = a \cdot E[\boldsymbol{s}_n], \tag{9}$$

in which the parameter a represents the constant increase in rate per unit time with dimension trials/s^2.
If $E[\boldsymbol{\lambda}_n]$ increases linearly over a length of time $\boldsymbol{s}_n$-$\boldsymbol{s}_0$, the average of $E[\boldsymbol{\lambda}_0]..E[\boldsymbol{\lambda}_n]$ must be equal to half the rate at time $\boldsymbol{s}_n$, i.e.:

$$\langle E[\lambda] \rangle_n = \frac{1}{n+1} \sum_0^n E[\boldsymbol{\lambda}_i] = \tfrac{1}{2} \cdot E[\boldsymbol{\lambda}_n] = \tfrac{1}{2} \cdot a \cdot E[\boldsymbol{s}_n], \tag{10}$$

in which $<E[\boldsymbol{\lambda}]>_n$ is the average expected rate over n trials, and with $s_0=E[\boldsymbol{\lambda}_0]=0$. It may be more insightful to consider the average rate in equation 10 may as the average speed of a car that travels with constant acceleration, i.e. with a linear increasing speed, starting with zero speed at time $s_0=0$. At any time $\boldsymbol{s}_n$, the average speed of the car over the elapsed travel time is half its momentary speed (rate), hence $<E[\boldsymbol{\lambda}]>_n = ½ \cdot E[\boldsymbol{\lambda}_n]$.
Also, average rate $<\lambda>_n$ over a length of time s_n-s_0 is per definition:

$$\langle \lambda \rangle_n = \frac{n}{s_n}, \quad (with\, n \geq 1), \tag{11}$$

and thus,

$$\langle E[\lambda] \rangle_n = \frac{n}{E[\boldsymbol{s}_n]} \tag{12}$$

By combining equations 10 and 12, we obtain an expression for the corrected practice time $\boldsymbol{s}_n$ that is required for n trials:

$$E[\boldsymbol{s}_n] = \sqrt{\frac{2}{a}} \sqrt{n}\,. \tag{13}$$

Since corrected practice time $\boldsymbol{s}_n$ equals actual practice time $\boldsymbol{t}_n$ minus the dead time in n trials, that is, $n \cdot T_d$, the expected practice time $\boldsymbol{t}_n$ for n trials is:

$$E[\boldsymbol{t}_n] = n \cdot T_d + \sqrt{\frac{2}{a}} \sqrt{n}\ . \tag{14}$$

To obtain an expression for the expected value of trial time T_n, we combine equation 14 and equation 6 to obtain:

$$E[\boldsymbol{T}_n] = E[\boldsymbol{t}_n] - E[\boldsymbol{t}_{n-1}] = T_d + \sqrt{\frac{2}{a}} \left(\sqrt{n} - \sqrt{n-1} \right). \tag{15}$$

Using a Taylor approximation for the part $\sqrt{n}-\sqrt{n-1}$, equation 15 can be rewritten as:

$$E\left[\boldsymbol{T}_n\right]=T_d+\sqrt{\frac{1}{2a}}\cdot\left(n-\tfrac{1}{2}\right)^{-\frac{1}{2}}. \qquad (16)$$

It follows from equation 16 that, with linear increase of performance in time, that trial times follow a power function with exponent –½ as a function of trial number *n*.
An expression for *a* follows directly from equation 14:

$$a=\frac{2n}{\left(E\left[\boldsymbol{t}_n\right]-n\cdot T_d\right)^2}. \qquad (17)$$

We point to an important property of the linear rate model. If we substitute the expression for parameter *a* (equation 17) in equation 9, we obtain:

$$E\left[\boldsymbol{T}_n\right]-T_d=\frac{E\left[\boldsymbol{t}_n\right]-n\cdot T_d}{2n}=\frac{\langle\boldsymbol{T}\rangle_n-T_d}{2}, \qquad (18)$$

in which $\langle\boldsymbol{T}\rangle_n$ denotes the mean over *n* trial times. This result illustrates that the expected value of each subsequent trial time is simply half the progressive average (i.e. the average of all preceding trial times), except for a constant T_d. Considering that the power law of equation 16 is equivalent to the 'progressive average function' of equation 18, it is comprehensible that averaging of trial times over trials or trainees causes a 'power law artifact', as argued by Anderson & Tweney (1997), Myung, Kim & Pitt (2000), and Heathcote, Brown & Mewhort (2000). However, if a series of non-averaged trial times from an individual can be described by an progressive average function, then all preceding trials are equally important for the success of a subsequent trial, irrespective of their individual duration or rank number in the process.
When we assume a quadratic increase in rate per unit time in equation 9, a power law can be derived with exponent $\alpha={}^2/_3$. Moreover, when we assume an exponential increase in rate of performance per unit time, the resulting power law will be hyperbolic, i.e. $\alpha=1$, corresponding with the power law that was generally found for noise free learning of a perceptron.
To summarize this section, we demonstrated that a power law with exponent $\alpha=½$ as a function of trials (examples) corresponds with a linear increase in performance (rate of trial completion) as a function of practice time when the task is speed-based. We also demonstrated that such power law is equivalent to a progressive average function.

2.6 The random trial-to-trial error and the test of models

Irrespective of the 'real' shape of the learning curve for complex tasks such as Space Fortress, a fitted smooth function for the expected values of subsequent trial times, be it a power function or some other function, does not account for apparently random trial-to-trial errors.
When contrasting the power law with other functions, such as the progressive average function

or the exponential function, we need to consider how the pioneers of the power law (Snoddy, 1926, de Jong, 1957, Newell & Rosenbloom, 1981) dealt with these errors when applying the law to trial time data. As a matter of fact these researchers treated the power law as a 'log-log-linear law' (Newell & Rosenbloom, 1981), which means that trial times obeying this law should be approximately on a straight line when plotted against trial number on double-logarithmic paper. Subsequently, these researchers applied linear least-squares regression in the double-logarithmic plane, sometimes in combination with visual judgment of goodness-of-fit. The implicit assumption in this procedure is the presence of stable additive Gaussian residual error in the double-logarithmic plane. Thus, when omitting Newell & Rosenbloom's n_0 parameter, these researchers assumed that the power law model of equation 2, is a noise model of the form:

$$\boldsymbol{T}_n = T_d + b \cdot n^{-\alpha} \cdot \boldsymbol{\varepsilon}_n , \qquad (19)$$

which, after a logarithmic transformation (in this case the natural logarithm with base e), gives:

$$\ln\left(\boldsymbol{T}_n - T_d\right) = \ln\left(b\right) - \alpha \ln\left(n\right) + \ln\left(\boldsymbol{\varepsilon}_n\right), \qquad (20)$$

in which the additive error term $\ln(\boldsymbol{\varepsilon}_n)$ was assumed to possess a Gaussian distribution. According to the implicit assumptions in this model, trial times $\boldsymbol{T}_n$ must be considered as the result of a stochastic process $\boldsymbol{\varepsilon}_n$, which is controlled by a deterministic trend $b \cdot n^{-\alpha}$, instead of the result of a deterministic process to which stochastic noise has been added. With respect to the curve fitting method in the double logarithmic-plane, it must be noted that the best fitting estimates for α and b, depend on the nature of $\boldsymbol{\varepsilon}_n$, which leads to arbitrary results when no additional assumptions are made.

In this study we make the assumption that the random trial-to-trial error $\boldsymbol{\varepsilon}_n$ has a log-Gaussian distribution, which is the distribution followed by the exponential of a Gaussian distributed random variable. This distribution arises when many independent random variables are combined in a multiplicative fashion. Although these assumptions justify the log-log-linear law, the log-Gaussian distribution needs to be verified with experimental data. Moreover, we could assume a theoretical value for α on the basis of the result of Györgyi & Tishby (equation 4) or Kabashima & Shinomoto (equation 5). When we additionally make an independent estimate for the dead time T_d (to be clarified hereafter), an estimate for $\ln(b)$ (and hence for b) follows from equation 20:

$$\ln(\hat{b}) = \frac{\sum_{i=1}^{n} \ln(T_i - \hat{T}_d)}{n} + \alpha \cdot \frac{\sum_{i=1}^{n} \ln(i)}{n} . \qquad (21)$$

We consider dead time T_d to be a constant parameter that locates the lower bound of the distribution of trial times $\boldsymbol{T}_n$. We can estimate T_d for each individual trainee by determining the lower bound of the empirical trial time distribution, independent of the remainder of the model to be fitted. We present the numerical method for estimating T_d in appendix A. When we compare the fit of the various models to measured trial times (in the Results section), the estimates for T_d will be the same across the models, so that goodness-of-fit differences between the models are independent of errors in the estimates of T_d.

The essence of the research is the comparison of different learning models to explain learning of complex tasks. More specifically, we fit power laws with exponent $\alpha=1$ (the general solution for perceptron learning of simple noise free tasks), a power law with exponent $\alpha=½$ (the solution for perceptron learning of Györgyi & Tishby when stochastic noise in the task is maximal) and a power law with exponent $\alpha={}^2/_3$ (Kabashima & Shinomoto's solution for perceptron learning when stochastic noise is maximal). We will assume that prior practice parameter n_0 (equation 2) is zero and we estimate the scale parameter b from the series of trial times of six trainees, using equation 21. We will also fit a function of the form:

$$\boldsymbol{T}_n = T_d + c \cdot e^{-\varphi \cdot n} \cdot \boldsymbol{\varepsilon}_n, \qquad (22)$$

in which the expected trial time is an exponential function of trial number n. According to Heathcote et al., an exponential function fits better than a power function in many data sets. We thus estimate the free parameters c and φ at the same assumptions that we use for the power laws, that is, a constant T_d and log-Gausssian error ε_n. It is straightforward to obtain estimates for c and φ at these assumptions with linear regression when applying a logarithmic transformation (as in equation 20).

In addition, we fit a recursive form of the progressive average function (equation 18) to the individual trial time data, which is:

$$\boldsymbol{T}_{n+1} = T_d + \tfrac{1}{2}\left(\langle \boldsymbol{T} \rangle_n - T_d\right)\boldsymbol{\varepsilon}_n, \qquad (23)$$

which has only one parameter (dead time T_d) to be estimated from the data, as outlined in appendix A. This model needs an initial value for the progressive average $\langle \boldsymbol{T} \rangle_n$, for which we choose the first trial time T_1 such that the fit start with T_2.

3 Method

3.1 Task

The task that was used was the Space Fortress (SF) game (Donchin, 1989), a laboratory PC-game, specifically defined and designed for training research. The SF task, displayed on a VGA monitor, consists of a fortress in the center of the screen surrounded by two concentric hexagons and a spaceship. The trainee is required to control the spaceship's flight in two dimensions with a joystick and to destroy the fortress as quickly as possible by shooting missiles from his ship using the fire button on top of the joystick. Simultaneously, the trainee must avoid being destroyed by the fortress or by mines. The mines, moving towards the space ship, appeared at intervals of four seconds. Trainees also had to distinguish between two types of mines, and react accordingly. The more difficult mine, can be identified by a letter that appears in the information panel at the bottom of the screen (prior to each five minute block of play, the trainee is presented with a new set of three letters that are used to identify 'difficult' mines). Appearance of a difficult mine requires the trainee to press the right ('identification') button on the mouse twice with an interval of 250-400 ms before the mine can be destroyed by a missile. The 'easy' mine can simply be destroyed by hitting it with a missile without pressing the identification button. However, falsely pressing the identification button for an easy mine makes this mine invulnerable for missiles, such that it cannot be eliminated and will either hit the ship or automatically disappear after 10 seconds. Since missiles fired at the fortress have no effect when a mine is present, the trainee can choose whether to avoid the invulnerable mine and wait for it to disappear or let it damage the ship. Another complication in this task is that the supply of missiles is limited, and the stock has to be monitored in the information panel at the bottom of the screen. Extra supply can be obtained by using 'resource opportunities'. The availability of these opportunities are indicated by a random sequence of symbols (&, #, $, %, !, etc.) which appear in the center of the display (beneath the fortress). When the $ symbol appears for the second time in a row, the trainee can get extra missiles by clicking the middle button of the mouse. The top-level strategy of this task essentially requires speed-based performance ('destroy the fortress as quickly as possible'). One trial on this task was defined as a destruction of the space fortress. With this task and this definition, time-on-task corresponded to the sum of trial times.

3.2 Participants

Six male university undergraduate students aged between 20 and 23, with normal vision, participated in the study. The trainees were recruited via an advertisement in the University magazine of Utrecht University. Initially 36 trainees were selected from a larger group of 51 candidates by means of the Aiming Screening Task, a task that is known to be a reasonable predictor for training success on this task (see Foss, Fabiani, Mane & Donchin, 1989). Median and modal score of the larger group of 51 candidates was 860 points, with a range of 600-1160 pts. A minimum aiming screening score of 740 points was required to participate in the study. As the current study is part of a larger training study (to be reported elsewhere), the six trainees in the current study are a balanced subset of the full set of 36 trainees who participated in the larger study. The six trainees selected for this study had obtained either 880 (3 trainees) or 860 points (3 trainees) on this screening task, with the same average screening score (870 points) as the full set of 36 trainees. The goal of this selection was twofold. A first goal was to ensure that the capabilities of all six trainees would be sufficient to demonstrate considerable progress during sixteen hours of training. A second goal was to obtain a coherent group, comparable in performance to control groups in earlier experiments (such as the groups ran by Foss et al, 1989, and Gopher, Weil & Siegel, 1989).

None of the trainees reported playing video games for more than 4 hours per week. Trainees were paid an equivalent of 220 Dollars (245 Euro), plus a bonus of 60 Dollars (68 Euro) upon completion of the experiment.

3.3 Procedure

The trainees received 16 hours of practice training as follows: Eight training days were spread evenly over a 5 week-period. On each day of training, the trainees practiced simultaneously in separate cubicles. Each day the trainees received three training sessions of 8 games of 5 minutes each, separated by two breaks of 20 minutes. In each game of five minutes the trainee performed as many task trials (destroying the fortress) as possible. A possibly uncompleted task trial at the end of a game was carried over to the next game. The effective time-on-task (cumulative trial times) of a session was thus 40 minutes. Trainees were allowed to take one-minute breaks between games; during those breaks they could study their game score over the last game and previous games. The game scores were presented automatically on the screen in the form of a graphical learning curve, i.e. game score plotted as function of game number.

3.4 Software and Equipment

The experiment room contained 6 individual computer stations in separate cubicles. Each computer station was equipped with an IBM PC and a joystick of type FlightStick (CH-products). The joysticks were modified for connection to an A/D converter card (DataTranslation) in the PCs. The fire-button on the joystick and three other response buttons were connected to a timer card in the PC. Each cubicle contained a camera system, which the experimenter used to monitor the course of the experiment.

The SF software was made available by the dept. of Psychology, University of Illinois at Urbana-Champaign. The software was modified to record additional parameters, the most important of which were: total time-on-task and trial times, with a timing accuracy of 50 milliseconds. Further software provisions were made to record trial times associated with uncompleted trials at the end of five-minute games and to carry these over to the next game.

3.5 Further Training Materials

After their screening test and well before the start of the experiment, the trainees received the instruction booklet for the SF game by mail at their home address. The instruction booklet specified the rules of the game and explained how to control the space ship. No reference was made to specific tactics or strategies. The trainees were instructed to study the booklet carefully before the start of the experiment.

4 Results

4.1 Relation between trial time and other performance measures

We use trial times (in this case, the times needed to destroy the space fortress) or its reciprocal (rate with which fortresses are destroyed) as a performance measure in the remainder of this analysis. Therefore we verified whether these are suitable measures of performance on complex tasks like the Space Fortress (SF) game. Usually, performance on SF is evaluated on the basis of ‘total game score’, that is, total number of points obtained in five minutes of practice. We refer to Mane & Donchin (1989) for a description of the events affecting total game score on SF. To evaluate the suitability of trial times we calculated the correlation coefficients between the full series of total game scores of each trainee and the number of completed trials (number of fortress destroyed) per game. This was done on the assumption that a high positive correlation between the two measures would indicate that both were equally valid as performance measures. The averaged correlation coefficient over all trainees was 0.977, with the lowest correlation being for Trainee 6 ($r=0.962$) and the highest for Trainee 2 ($r=0.994$).

We also verified whether the total game score of our group was comparable to that of similar control groups run by Foss et al. (1989) at the University of Illinois and by Gopher et al. (1989) at the Technion Institute of Technology, Israel. These groups were trained under comparable circumstances, i.e., whole-task training without the use of any particular instruction technique.

Foss et al. report an average of 3280 pts for a group of 40 subjects after 10 hours of practice time and Gopher et al. report 2061 pts for a group of 24 subjects, also after 10 hours of practice time. The average game score of our group after 10 hours of training was 2611 pts (with a standard deviation of 1125 pts).

4.2 Analysis of empirical learning curves

Figure 1 shows learning curves as a function of time for each of the trainees. Trainees have been ranked according to the total number of trials completed during training, Trainee 1 having performed the highest number of trials. In the upper panel, performance is plotted as the number of trials (number of fortresses destroyed) per session of 40 minutes. The plot in the lower panel is based on cumulative numbers of trials, i.e. the total number of trials that have been completed at t, against t. At each trial, the cumulative plot jumps one upwards. From both curves it is clear that all trainees demonstrated a global performance improvement and the combination of both types of curves gives an impression of the differences in the ability of the trainees. More specifically, a number of individual differences can be observed (summarized in table 1):

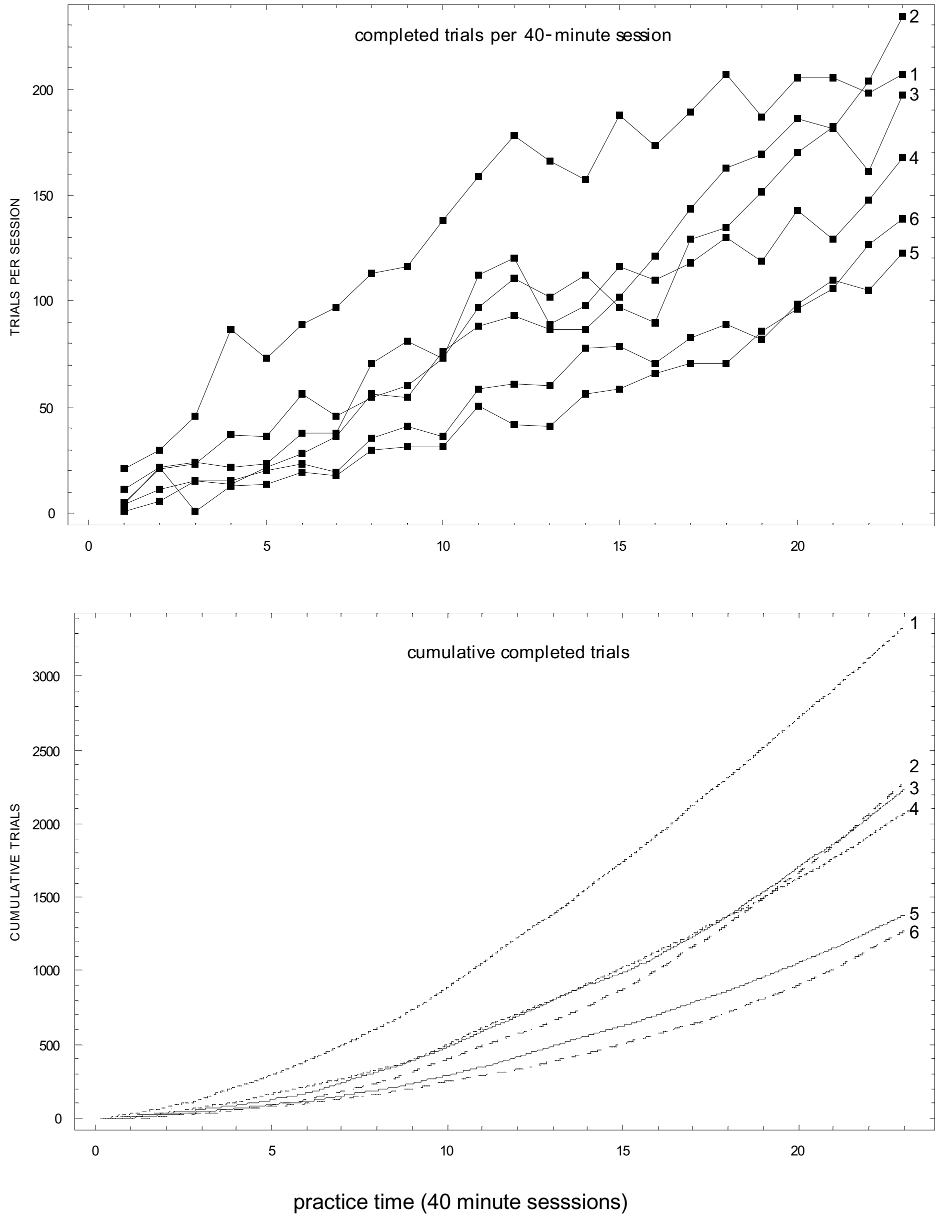

Figure 1: Individual learning curves for each of the six trainees (numbered 1 to 6) – Completed trials per 40 minute session (above) and cumulative completed trials (below).

Table 1

Summary of global individual characteristics of trial time data for 6 trainees – Column 1: trainee number, trainees are sorted according to the total number of trials completed (Column 2, see also figure 1). Column 3: Maximum observed trial time. Column 4: Minimum observed trial time. Column 5: Mean of all trial times. Column 6: Standard deviation in all trial times. Column 7: Mean increase in performance (completed trials per session) per session.

Trainee	Total trials	Max. trial time [s]	Min. trial time [s]	Mean trial time [s]	Standard Deviation Trial times [s]	Mean improvement $\left[\frac{trials}{(session)^2}\right]$
1	3346	638	5.8	16.2	16.4	8.2
2	2282	1399	4.6	22.4	47.9	9.7
3	2229	853	6.1	23.3	31.7	8.0
4	2090	1098	5.8	24.9	30.9	6.8
5	1379	867	7.7	35.5	42.0	5.0
6	1273	1433	7.3	38.5	66.3	5.3
Mean	2175	1048	6.2	26.8	36.5	7.2

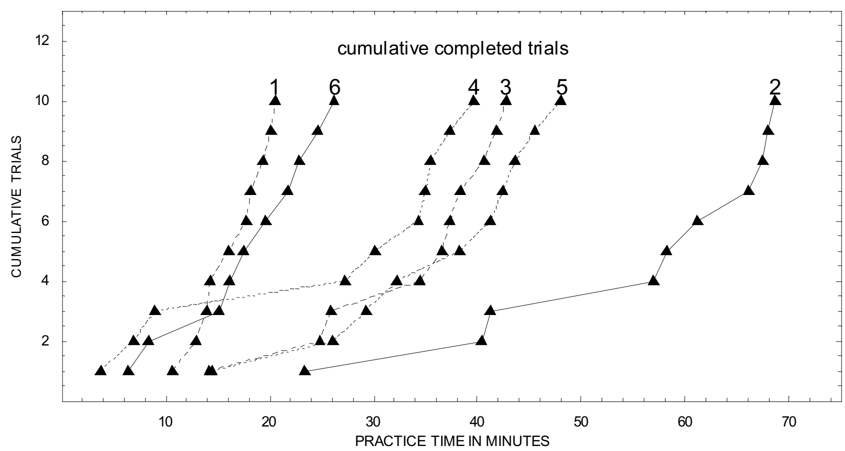

Figure 2: Cumulative plots for the first ten trials of each of the six trainees (numbered 1 to 6).

– The curves that plot performance against time in the upper panel of figure 1 all seem to be different. Some of these curves seem to decelerate (Trainee 1); others seem more linear or seem to accelerate (Trainee 2). Some curves have local sections where improvement stops temporarily or drops.

- There are vast differences in training performance, when measured according to the number of total trials performed. The poorest performer in this respect (Trainee 6) completed less than 40% of the number of trials completed by the best performer (Trainee 1).
- Also, considerable differences were observed in improvement per unit time, expressed as the slope of the straight line connecting the score (trials per session) in the first session with the score in the last session. The rate of improvement of Trainee 5 was almost twice as low as that of Trainee 2.
- Initial performance was generally unstable, and initial trial times fluctuated considerably. Individual differences at the start of the training are represented in figure 2, in which cumulative trials are plotted against time for the first ten trials. The plot shows that Trainee 2 took well over an hour to complete the first ten trials, while Trainee 1 needed only 20 minutes. Although Trainee 2 was initially the slowest, in a later stage he clearly outperformed other trainees with respect to speed. The plot generally demonstrates that initial training performance as measured by these curves is not representative for later performance.
- It is evident from both the maximum and minimum observed trial times that trainees are sometimes capable of completing trials in 1% of the time that they initially needed. This wide range of individual trial times indicates large individual variations, which is also expressed in table 1 by means of the standard deviation.

When trial times are plotted against trial-number rather than numbers of aggregated trials against time (figure 1), a different picture emerges (figure 3). The large dispersion in trial times is now apparent. We only depict the raw data for Trainee 1 and Trainee 6, but the pictures are typical for all trainees. Both the average and the standard deviation of trial times decrease as a function of trial number. Unlike the time curves of figure 1, the curve of figure 3 reveals that improvement as a function of trials slows down: the average trend is concave and eventually flattens out. A lower limit to the trial times is clearly visible: Trainee 1 has no trial times shorter than 5.8 seconds and trainee 6 has no trial times shorter than 7.3 seconds (see table 1), but values approaching this lower limit are observed rather early in the process. A constant 'dead time', not subject to learning, seems to determine the absolute minimum of trial times, which forms a constant limitation to performance. This dead time arises partly from rules pre-programmed in the Space Fortress game: intervals between missile shots fired from the space ship need to last at least 250 ms, otherwise the trainee is penalized. Since at least ten hits are needed to destruct the Fortress, at least 2.5 s of dead time in each trial time is included in the task (for details see Mané & Donchin, 1989, p.18). Dead time may be considered as the sum of the constant delays imposed by the task, not only the delays pre-programmed in the task, but also the constant 'incompressible' delays attributable to the trainee.
Two more features of figure 3 are notable with respect to the probability distribution of trial times; First, short trial times are more likely than long trial times, a tendency that is stronger towards the end of the curve, and second, extreme trial times are more likely to be extremely positive than to be extremely negative with respect to the mean. The asymmetry is a fundamental consequence of the fact that trial times are bounded below but are unbounded

above. Thus, trial times have a probability distribution that is skewed towards small values and with a long high tail.

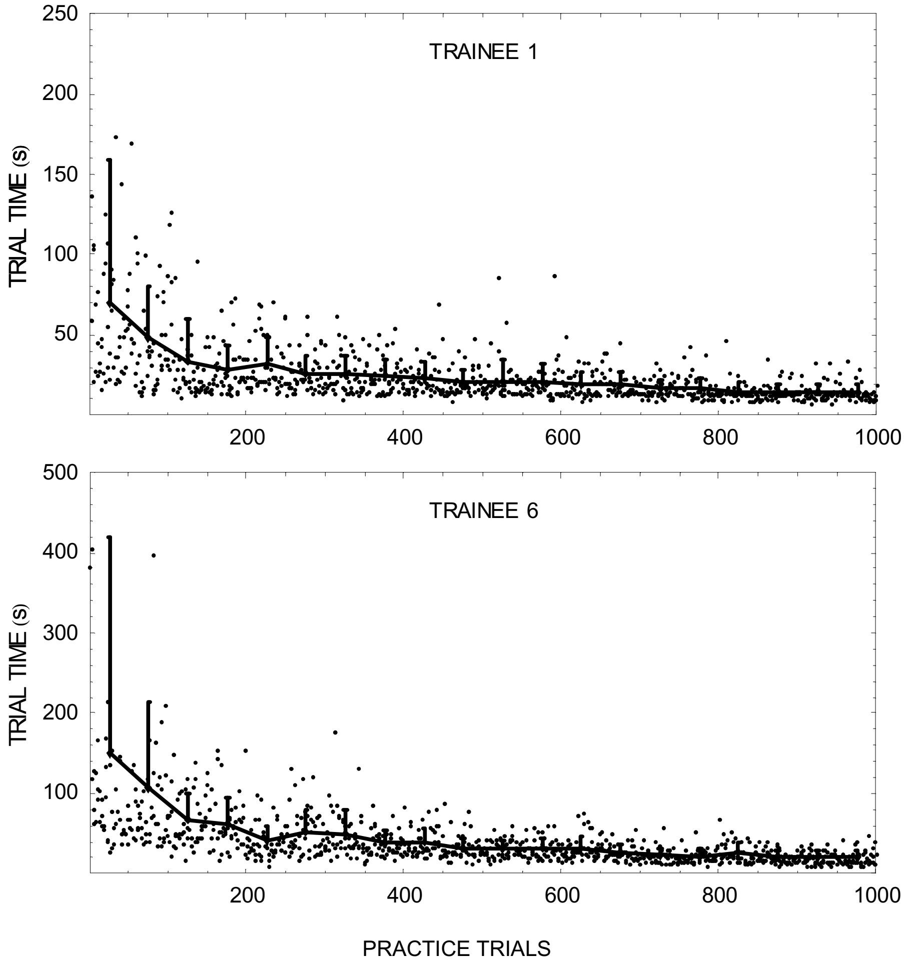

Figure 3: The first thousand trial times plotted against trial number for Trainee 1 and Trainee 6. The line connects average trial times of subsequent series of 50 trials. The error bars denote the standard deviation in each series. Note the different time scales for the vertical axes.

4.3 Model fits

We examine the fits of five possible learning curve functions to the individual series of trial times, such as these depicted in figure 3, using the fitting methods as described in the theory section. First, the power function of equation 19 with exponent α=1 (a hyperbolic function); Second, the exponential function of equation 22, proposed by Heathcote, Brown & Mewhort (2000); Third, the power function with exponent $\alpha={}^2/_3$; Fourth, the power function with exponent α=½; Fifth, the progressive average function of equation 23. Hence, these fits are based on estimation of the dead time parameter T_d as outlined in appendix A. For all other free parameters, the estimates are maximum likelihood estimates at the assumption of multiplicative log-Gaussian error, which will be verified later.

In figure 4a we represent the best fitting hyperbolic function and the best fitting exponential function, together with a moving average of the raw data (merely to visually evaluate goodness-of-fit). In figure 4b we represent the best fitting power laws with exponents $\alpha={}^{2}/_{3}$ and $\alpha=½$, respectively, and the progressive average function (with one free parameter). To compare the fits of the different models graphed in figure 4a and 4b, we will use four different statistics: MPE, MAPE, RMSE and R^2. The first three measures are error-measures and indicate lack-of-fit. The latter is the squared correlation between the trial times predicted by the model and those observed, and thus indicates goodness-of-fit. Neither of these four measures alone is sufficient to decide on the correctness of a model. However, in considering a combination of different goodness-of-fit measures, we may be able to appreciate the differences between the models. We shortly explain the four statistics.

Mean Percentage Error (MPE, StatSoft, 2001). A purpose of the presented models is to describe or predict performance of a trainee, for example in terms of the trend in trial times over some period, or, equivalently, in terms of trials or practice time required until the trend in the curve exceeds a certain criterion skill level. In both examples, the most appropriate lack-of-fit measure is the sum of residual errors in the trial times (i.e. predicted trial time minus observed trial time, summed over all trials). The sum of errors provides us with the gross error in predicted practice time. To obtain a general reference, we compute the relative sum of errors by dividing the sum of errors by the sum of observed trial times. Multiplication by a factor hundred yields the Mean Percentage Error (MPE). A characteristic of the MPE is that positive and negative error values in individual trial times neutralize each other, such that this measure is a poor indicator of overall fit. Thus, a low MPE is a condition for a good fit, but not a guarantee for a good overall fit.

Mean Absolute Percentage Error (MAPE, StatSoft, 2001) The MAPE is computed as the average absolute error value, expressed as a percentage. The MAPE reflects that the average prediction of a single trial time by the model is "off" by some percentage. The latter can be useful information. However, when the distribution of errors is asymmetrical, the MAPE will not be minimal, even if the learning curve describes the expected values.

Root of Mean of Squared Errors (RMSE) and R^2 The commonly used RMSE and the percentage of variance accounted for (R^2) are both based on squared errors and therefore emphasize extreme error values. These measures are optimal when the errors are additive and have a Gaussian distribution, which conditions are clearly not fulfilled.

We summarize those measures, averaged over the six trainees, in table 2. A complete list of parameter values and goodness-of-fit measures for each model and for each individual trainee is included in appendix B.

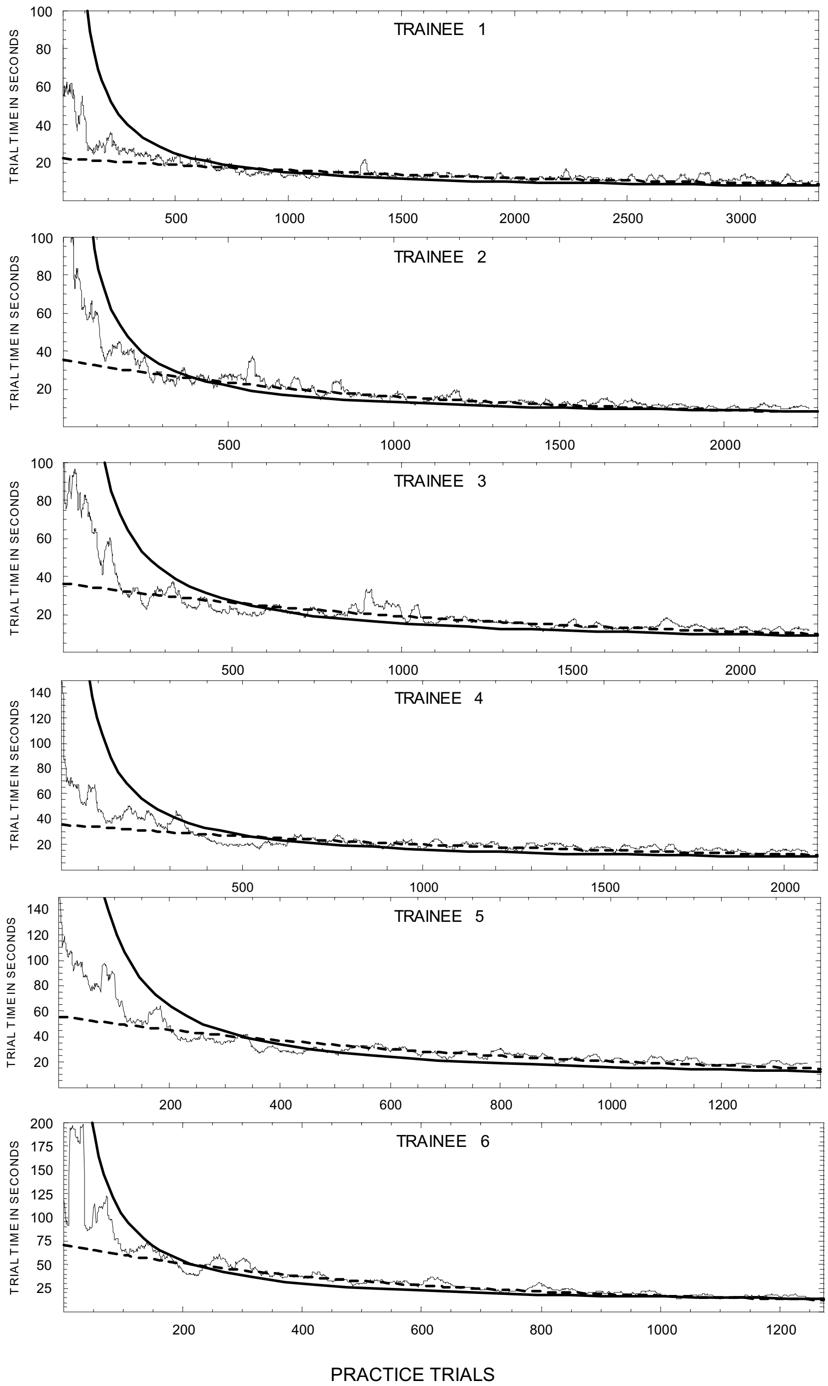

Figure 4a: Fits of the hyperbolic function (solid lines) and exponential function (dashed lines). The jagged solid lines are the moving averages of the raw data (window of 25 trials).

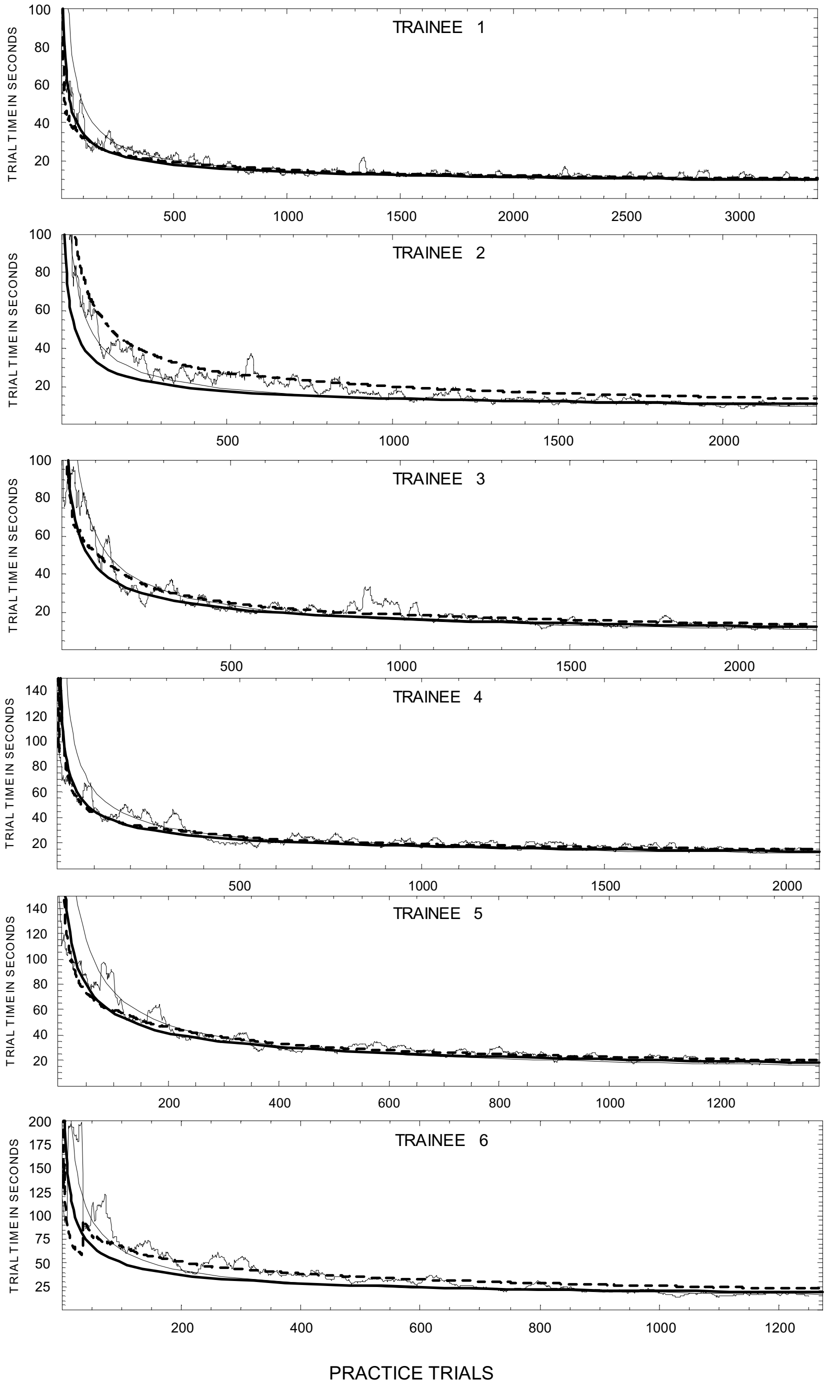

Figure 4b: Fits of the progressive average function (thick dashed lines), power law with exponent α=½ (thick solid lines) and power law with exponent α=$^2/_3$ (thin solid line). The jagged solid lines are the moving averages of the raw data (window of 25 trials).

Both figure 4a and table 2 show that the hyperbolic function ($\alpha=1$) provides an inadequate description of the trend in the data. The function predicts too large trial times at the start of the curve and reduces too fast towards the asymptotic trial time.
The exponential function also provides an inadequate description of the trend in the data. The curve importantly underestimates the trend in trial times for the first few hundred trials.
The hyperbolic function and the exponential function are inferior to the other three models in all six cases (the fit statistics, for each trainee separately, are listed in appendix B).
Figure 4b shows that the distinction between the power functions with $\alpha={}^{2}/_{3}$ and $\alpha=½$ with respect to goodness-of-fit is not large. However, the MPE for the ${}^{2}/_{3}$-curve is better. This is a practical and important measure when the model must predict total practice time or average trial time. In some cases (Trainee 2 and Trainee 6), the ½-curve substantially under-estimates the trend in the trial times at the beginning of the curve, where large deviations from the average trend are more common. In these cases, the ${}^{2}/_{3}$-curve seems to provide a better fit.

Table 2
Summary of fits for five different learning curve models. Fit statistics have been averaged over six trainees. From left to right: Mean Percentage Error (MPE) in the models' estimation of trial times. Mean Absolute Percentage Error (MAPE) and Root Mean Square Error (RMSE) in the models' estimation of trial times. Percentage values for the models' coefficient of determination R^2.

		FIT STATISTICS			
	Model	**<MPE>** [%]	**<MAPE>** [%]	**<RMSE>** [s]	**100*<R²>** [%]
1	$E[\boldsymbol{T}_n] = \hat{T}_d + \hat{b} \cdot n^{-1}$	86.3	61.8	283.1	43.9
2	$E[\boldsymbol{T}_n] = \hat{T}_d + \hat{c} \cdot e^{-\hat{\varphi} \cdot n}$	-17.8	32.9	36.7	17.9
3	$E[\boldsymbol{T}_n] = \hat{T}_d + \hat{b}' \cdot n^{-\frac{2}{3}}$	-1.1	34.9	37.0	48.9
4	$E[\boldsymbol{T}_n] = \hat{T}_d + \hat{b}'' \cdot n^{-0.5}$	-14.9	31.9	30.7	47.8
5	$E[\boldsymbol{T}_{n+1}] = \frac{1}{2}\hat{T}_d + \frac{1}{2}\langle \boldsymbol{T} \rangle_n$	-0.1	41.1	28.1	51.0

The recursive progressive average function, based on the progressive average $\langle \boldsymbol{T} \rangle_n$, avoids substantial underestimation. In the absence of random residual error, the recursive model would be equivalent to the power law with $\alpha=½$. However, in the presence of random error the progressive average function provides a better description of the trend in the data than the corresponding power law (and the other power laws) as measured by the MPE, RMSE and R^2.
The MAPE of the progressive average function is larger than for the power laws with $\alpha={}^{2}/_{3}$ and $\alpha=½$. The MAPE indicates that the average prediction of a single trial time by the recursive progressive average function is "off" by 41.1 percent. This larger MAPE does not necessarily imply an inferior fit when comparing different models, because it depends importantly on

skewness of the error distribution. For example, the theoretical MAPE that would result if the errors follow the moderately skewed Weibull distribution (with shape parameter 2) is 42%. If the errors were distributed according to the skewer negative exponential distribution (or, likewise, a Weibull with shape parameter 1), the MAPE would be larger, approximately 74%.
On average, the recursive progressive average function fits better than any of the power laws. However, in the individual case of Trainee 2, the progressive average function clearly overestimates the trend. In this case, the power law with $\alpha={}^2/_3$ seems to fit the data better, but neither of the five models is satisfactory in terms of MPE (see Appendix B), due to persisting irregularities in the learning curve of Trainee 2.

4.4 The distribution of trial times

In the previous section we made the assumption that the random trial-to-trial errors are multiplicative and are distributed according to a log-Gaussian distribution. We verified with the correlogram-method (Cox & Lewis, 1966, p.93) whether the random error variables E can be considered independent and coming from an identical distribution. The errors E that resulted from the fit with the exponential function and the power law with α=1 revealed substantial correlation, indicating that these functions did not capture the trend in the data (which is also clear from figure 4a). Since the recursive progressive average function yielded the best fits, the multiplicative random error E for each trial is best approximated by using equation 18, giving:

$$\hat{\varepsilon}_n = \frac{2(T_n - \hat{T}_d)}{\langle T \rangle_n - \hat{T}_d}. \qquad (24)$$

In figure 5, we present the histograms of these errors E for each of the six trainees in the experiment and for the full number of trials of each trainee. We superimposed the best fitting log-Gaussian density function on these histograms. The log-Gaussian parameters were estimated with the maximum likelihood method, as advised by van Zandt (2000). These histograms show that the distribution has a long stretched high tail, which is furnished by the log-Gaussian distribution.
To compare the empirical distributions with the best fitting log-Gaussian, we calculated the coefficients of variation (CV, the ratio of standard deviation and mean) in the error E for each trainee and compared these with the coefficients of variation for the best fitting log-Gaussian distribution. Both empirical and fitted CV's are approximately 0.7, with a substantial deviation for trainee 6, whose CV is 1.03.
However, it must be noted that the CV relies on the standard deviation, which is based on squared deviations and thus emphasizes extreme values. If the data contains a few outliers, it will predominantly affect the standard deviation, which is apparent in the data of Trainee 6. It is also notable that with more basic tasks considerable lower CV's have been reported, in the order of 0.2 to 0.4 (e.g. Logan, 1992).

Table 3
Distribution Parameters. From left to right, Mean error, Standard Deviation in errors (SD), and coefficient of variation (CV, the ratio of standard deviation and mean), CV of the best fitting log-Gaussian distribution, relative deviation of the log-Gaussian CV from the empirical distribution.

Trainee	Mean ε	SD ε	CV ε	CV Log-Gauss	Deviation [%]
1	1.02	0.72	0.71	0.69	-2.39
2	0.78	0.57	0.73	0.71	-2.11
3	0.88	0.59	0.66	0.68	2.81
4	1.03	0.66	0.64	0.71	10.77
5	0.94	0.56	0.60	0.61	2.05
6	0.90	0.93	1.03	0.77	-25.86
Mean	0.93	0.67	0.73	0.69	-2.46

A different way of characterizing the distribution of the error variable is by way of its hazard function. In the absence of parameter variability, the hazard function $h(\varepsilon)$ is more diagnostic of a particular distribution than the density function $f(\varepsilon)$ or the cumulative distribution $F(\varepsilon)$ (van Zandt, 2000). The hazard function (or hazard rate) h is defined as:

$$h(\varepsilon) = \frac{f(\varepsilon)}{1 - F(\varepsilon)}. \qquad (25)$$

Roughly speaking, the hazard rate of the random error variable ε can be interpreted as the probability of immediate completion of the trial, given that the increasing error has a value of ε already.

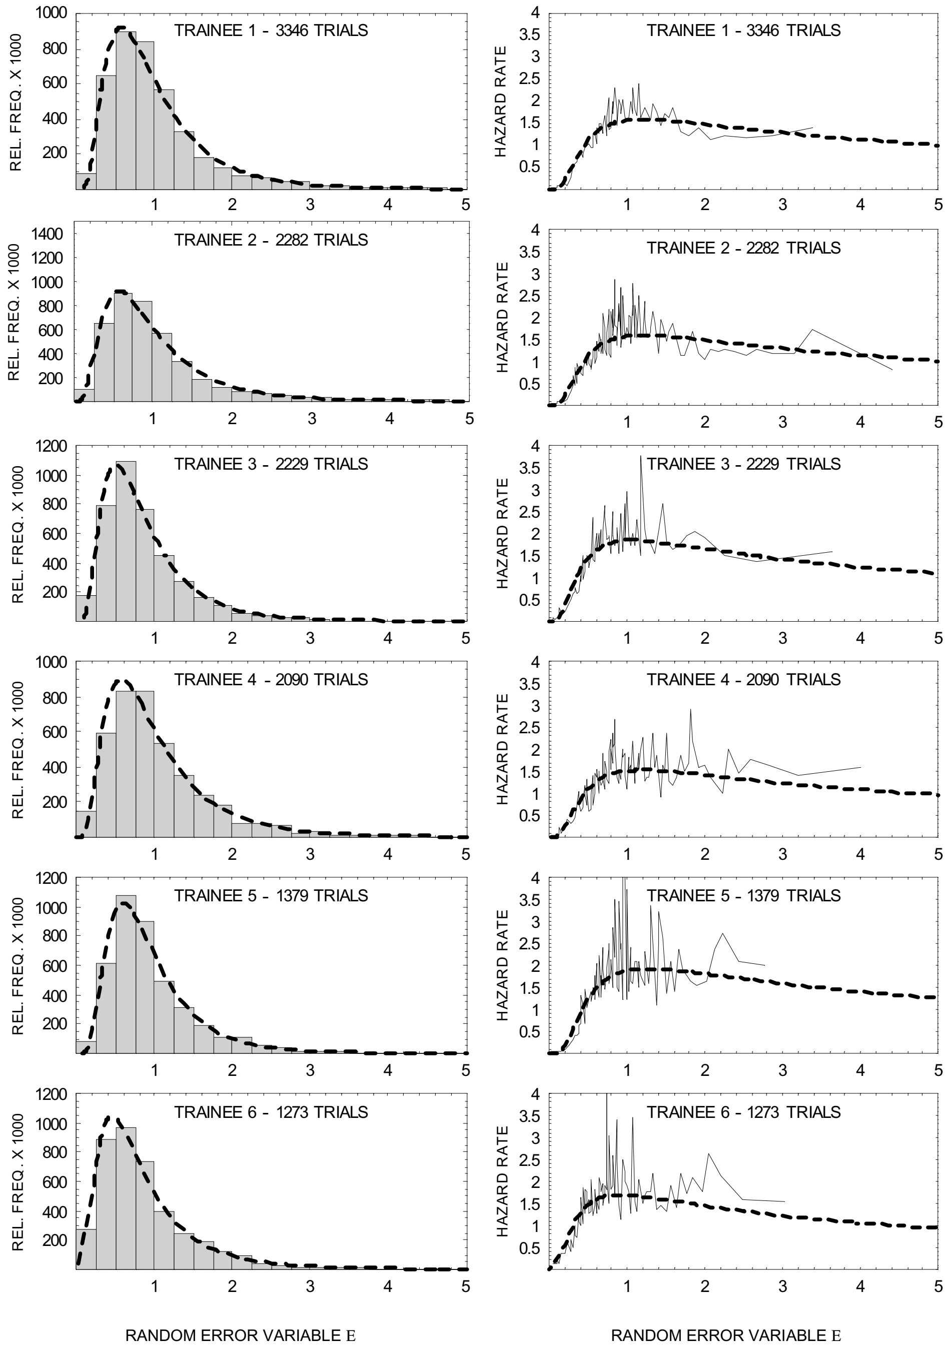

Figure 5: Empirical histograms and plots of the empirical hazard rate of the multiplicative error in trial times. The best fitting log-Gaussian density function (dashed line) is superimposed on the histograms. The corresponding log-Gaussian hazard function is superimposed on the plots of the hazard rate.

In figure 5 we present the empirical hazard functions, which have been computed according to the numerical method used by Luce (1986, p.123). Although the empirical hazard functions are noisy, they appear to be peaked functions, i.e. increasing to a maximum and subsequently decreasing. This shape of hazard function is indeed characteristic for a log-Gaussian distribution and excludes a number of possible alternatives, that is, other skewed distributions that have been proposed for modeling of response times. On the basis of figure 5, it is clear that trial times do not follow the commonly used *negative exponential distribution*, whose hazard rate is constant.

The *ex-Gaussian distribution*, which is the convolution of a normal distribution and a negative exponential distribution, often gives good fits to simple reaction times (van Zandt, 2000). However, its hazard function is an increasing function (Luce, 1986), such that it will not fit the empirical hazard functions of figure 5. Another alternative is a *Weibull distribution*, which has been proposed to model the distribution of response times in a learning process, when the trend follows a power law as a function of trials (Logan, 1988a, 1992). However, Weibull hazard functions are themselves power laws, i.e. monotonously increasing or monotonously decreasing functions, such that a Weibull hazard function will not fit the empirical hazard function.

A *Gamma distribution* (the convolution of a series of identical negative exponential distributions) has a hazard function that first increases and then levels off at a constant maximum value. At least in some cases the empirical hazard functions of figure 5 seem to decrease after reaching a maximum. In fact, the decreasing high tails of the empirical hazard functions of figure 5 signify that the probability of immediately finishing a trial decreases when the trial time becomes large.

5 Discussion of results

The empirical learning curve, that is, a measure of skill-level as a function of trials or time, is an important manifestation of skill acquisition. However, the empirical probability distribution of the performance measure is equally important. In case of a speed-based task, skill level can be measured by trial times, or its reciprocal, the rate of trial completion.
We measured series of trial times (the times needed to complete subsequent trials) during long practice training with a complex speed-based task, the Space Fortress game. By means of analysis of these time series, we attempted to model both the learning curve and the distribution of the trial times of individuals during the acquisition of complex skills. Learning a complex skill from practice must be considered as a stochastic process and the learning curve is merely a trace of *expected* skill level.

5.1 The linear rate model

The overall trend analysis with five competing models shows that a series of trial times is best modeled by the recursive progressive average function of equation 23. This function, which follows from the linear rate model, states that each subsequent trial time is expected to be half the average of all preceding trial times, plus half the dead time inherent in each trial. Thus, all preceding trials are equally important for the expected skill level at a subsequent trial, which denies any deterministic order effect in learning a complex task.
When there would be no trial-to-trial error, the progressive average function would correspond with a deterministic learning curve that is a power function of trials with (negative) exponent ½. However, we found that the power law with exponent ½ provides a worse fit to the data than the progressive average function, which can only be attributed to the peculiar nature of the random trial-to-trial fluctuations early in the learning process. Also, in the presence of such fluctuations, it is difficult to distinguish a power law with exponent ½ (Györgyi & Tishby, 1990) from a power law with exponent $^2/_3$ (e.g. Uezu & Kabashima, 1996). Moreover, a power law has an extra scale parameter (which need to be estimated) and is therefore a less parsimonious description of the data than the progressive average function.

5.2 The probability distribution of trial times

When we fitted a power law with a constrained exponent α to a series of trial times, we needed to estimate the free scale parameter b from the data. In this process, we made the explicit assumption that the trial-to-trial error was multiplicative and log-Gaussian. This corresponded with the implicit assumption of some researchers (foremost Newell and Rosenbloom, 1981) when fitting the logarithmic form of the power law (equation 20) to logarithmic transformed trial times with least-squares regression. In this case, the appropriateness of the method of least squares regression depends on the assumption that log-transformed trial times have additive error with a Gaussian distribution. The fits of these latter researchers were generally satisfactory in terms of R^2. However, the reported parameter estimates seemed arbitrary (due to their inseparability, as we argued). In spite of the widespread acceptance of the power law as a fundamental benchmark result and its tremendous influence on the development of theories of automaticity (Palmeri, 1999), the assumptions on the distribution of trial times were never verified. In this research, we exploited the progressive average function to recover the multiplicative trial-to-trial error. We found close fits of the multiplicative trial-to-trial error to the log-Gaussian distribution, with a typical value for the coefficient of variation (the ratio of standard deviation and mean) of 0.7.

5.3 A naïve cognitive strategy for learning

In the introduction we referred to learning of perceptrons when these are tasked with the prediction of a desired pattern from a noisy example, that is, when the target rule is difficult to extract. These learning machines demonstrate a rich range of behaviors in the first trials, but after a large number of trials they exhibit power law behavior.

Analogous to perceptron learning we can speculate on a cognitive mechanism for a human learner, which leads to the reduction in trial times, in accordance with the recursive progressive average function. Such a strategy could be based on the storage of patterns of perception and action in memory. These patterns represent exposures to the task. During a new exposure to the task, one such pattern, randomly retrieved from the available set of patterns, will guide this exposure in terms of perceptions and actions required for finishing the task. If the current exposure to the task is deemed more successful than the guiding pattern, the current exposure will create a new pattern, which may guide future exposures to the task. Alternatively, when the current exposure is unsuccessful relative to the guiding pattern (i.e. when the guiding pattern had better prospects than the actual outcome of the exposure), the guiding pattern will create an approximate copy of itself, which may also guide future exposures to the task. In this chain of events, a newly stored pattern will generally be better than an old pattern. More specifically, with random retrieval of one pattern, and creation of one new pattern on each exposure as the result of one comparison, the new pattern will be on average twice as successful as the patterns that guided previous exposures.

In a speed-based setting, successful patterns lead to faster trial times. When we consider each complete trial on the task as one exposure to the task, repeated comparison leads to trial times whose expected value is exactly half the progressive average trial time (except for the dead time term, i.e. conform equation 23).

In the introduction we mentioned that the essence of complex skill acquisition is to learn to select the most effective methods that will lead towards the target of the whole task. In this context we consider a method as a pattern of perceptions and actions. An accumulative learning mechanism based on the repeated comparison and storage of these patterns (comparing a new method against an existing, stored, method) may be a feasible cognitive strategy in difficult learning situations. This may be particularly true in settings where the relationship between the current state of the task and the goal state seems stochastic to the learner and no other guidance, such as on-line instruction, is available.

5.4 Stability with Practice

The evidence in this experiment, in which no particular instruction was given to the trainees, clearly points to learning based on accumulation. There are important differences between learning on the basis of accumulation and learning on the basis of replacement (see also Restle & Greeno, 1970). In Crossman's (1959) replacement model, the operator has a fixed repertoire of methods, and only summary statistics (strengths) need to be updated in memory to change the probabilities for selecting each method. In contrast, accumulation mechanisms require a large memory, in order to store each new method.

As is directly apparent from the progressive average function of equation 23, the proposed accumulation process becomes more resistant to change with each practice trial. In other words, the transfer process to a new set of methods to do the task is slow and difficult. We speculate that if an individual has accumulated a large number of trials in which he or she practiced with sub-optimal methods to execute the task, and later discovers more appropriate methods, the old methods can only slowly be 'averaged out' in favor of the more appropriate methods.

5.5 Instance theory

The speculated mechanism is not a unique explanation to account for the typical reduction in trial times. One of the most discussed theoretical models in learning literature is the one designed by Logan (1988a, 1992). Although the theory provides an explanation for the retrieval of declarative accumulated knowledge from memory, rather than an explanation for the broader process of executing a complex task, the theory suggests implications for the design of training programs and the understanding of expertise (Logan, 1988b).

Logan models learning from practice in terms of a race between instances (memory representations) and an algorithm, the latter being a rather slow series of steps that takes the learner through the task. As more instances accumulate with practice trials, it becomes more likely that one of these instances will win the race from the algorithm. Logan's instance theory, based on the accumulation of instances, accounts for a power law as a function of trials and, more generally, for the change of the entire trial time distribution with practice.

Logan (1992) assumes, by invoking a theorem from extreme value statistics, that memory retrieval times converge quickly to a Weibull distribution during a learning process, the process being modeled as a race between many instances. The combination of the gradual increase in the number of fast instances (the winners of the race) and the consequent Weibull shape of their retrieval time distribution would cause subsequent trial times to follow a power function with an exponent that is the inverse of the shape parameter of the Weibull distribution.

The finding of an (approximate) power law in combination with a longer tailed log-Gaussian distribution does not seem to agree with the predicted convergence to a Weibull distribution for the case of learning a complex tasks. However, convergence of a process towards a limiting distribution may be slow (see e.g. the comments of Colonius, 1995, on the instance theory) or maybe disturbed by other processes than learning (e.g. forgetting, motivational effects, etc.). This study indicates that in the process of learning a complex task, a power law is not exclusively bound up with a Weibull distribution of trial times.

6 Conclusions

The training experiment in which each trainee received 16 hours practice with a complex task provided evidence that series of trial times are valid and extremely informative indicators for individual skill acquisition. We quantified the statistical properties of these series of trial times. On a local scale, trial times are independent stochastic variables and come from a long-tailed distribution. On the basis of analysis of trial-to-trial fluctuations, we suggest that the distribution of trial times is approximately a log-Gaussian distribution. On the long-term, this stochastic process is controlled by a deterministic trend.

To model this trend, we propose a linear relationship between practice time and the rate of trial completion in complex task learning (the linear rate model). The progressive average function, which follows from the linear rate model, provides a more robust description of the average course of individual trial times than the regularly proposed power laws and exponential functions. The model predicts that each trial time is half the progressive average of the preceding trial times plus half the dead time. In the absence of trial-to-trial error this model corresponds with a power law with (negative) exponent ½. This non-trivial power law indicates a link between human learning of complex tasks and neural learning of the recognition of patterns in the presence of maximum noise in the training examples.

7 Acknowledgement

The authors thank Dr. John van Rooij (RNLA) for enabling the data collection.

8 References

Anderson, J.R. (1984). The architecture of cognition. Cambridge: Harvard University Press.

Anderson, R.B. & Tweney, R.D. (1997). Artifactual power curves in forgetting. Memory and Cognition, Vol. 25, 724-730.

Boff, K.R. & Lincoln, J.E. (1988). Engineering Data Compendium: Human Perception and Performance. Wright-Patterson AFB: AAMRL.

Colonius, H. (1995). The instance theory of automaticity: why the Weibull? Psychological review, Vol. 102, No. 4, 744-750.

Cox, D.R., Lewis, P.A.W. (1966). The statistical analysis of series of events. London: Methuen & Co.

Crossman, E.R.F.W. (1959). A theory of the acquisition of speed-skill. Ergonomics, Vol. 2, 153-166.

Donchin, E. (1989). The Learning Strategies Project, introductory remarks. Acta Psychologica, Vol. 71, 1-15.

Donchin, E., Fabiani, M. & Sanders, A. (1989). The Learning Strategies Project: an examination of the strategies in skill acquisition. Acta Psychologica, Vol. 71.

Engel, A. & Broeck, C. van den (2001). Statistical Mechanics of Learning. Cambridge University Press.

Fisher, R.A. & Tippett, L.H.C. (1928). Limiting forms of the frequency distribution of the largest or smallest member of a sample. Proceedings of the Cambridge Philosophical Society, Vol. 24, 180-191.

Foss, M.A., Fabiani, A., Mané, A.M. & Donchin, E. (1989). Unsupervised practice; the performance of the control group. Acta Psychologica, Vol. 71, 23-51.

Gallistel, C.R. & Gibbon, J. (2000). Time, rate and conditioning. Psychological review, Vol. 107, No. 2, 289-344.

Gopher, D., Weil, M. & Bareket, T. (1992). The transfer of skill from a computer game trainer to actual flight. In: Proceedings of the Human Factors Society, 36th annual meeting.

Gopher, D., Weil, M. & Bareket, T. (1994). Transfer of skill from a computer game trainer to flight. Human Factors, 36(3), 387-405.

Gopher, D., Weil, M. & Siegel, D. (1989). Practice under changing priorities: An approach to the training of complex skills. Acta Psychologica, Vol. 71, 147-177.

Györgyi, G., & Tishby, N. (1990). Statistical theory of learning a rule. In W.K. Theumann & R. Köberle (Eds), Proceedings of the STATPHYS 17 workshop on neural networks and spin glasses (pp. 3-36) Singapore: World Scientific.

Hart, S.G. & Battiste, V. (1992). Field test of video game trainer. In: Proceedings of the Human Factors Society, 36th Annual Meeting.

Heathcote, A., Brown, S. & Mewhort, D.J.K. (2000). The Power Law repealed: The case for an exponential law of practice, Psychonomic Bulletin & Review, Vol. 7, 2, 185-207.

Jong, J.R. de (1957). The effects of increasing skill on cycle time and its consequence for time standards. Ergonomics, Vol 1, 51-60.

Lee F.J. & Anderson, J.R. (2001). Does learning a complex task have to be complex?: A study in learning decomposition. Cognitive Psychology, Vol. 42, 267-316.

Logan, G.D. (1988a). Towards an instance theory of automatization. Psychological review, Vol. 95, No. 4, 492-527.

Logan, G.D. (1988b). Automaticity, resources, and memory: theoretical controversies and practical implications. Human Factors, Vol. 30, No. 5, 583-598.

Logan, G.D. (1992). Shapes of reaction-time distributions and shapes of learning curves: a test of the instance theory of automaticity. Journal of Experimental Psychology; Learning, Memory and Cognition, Vol. 18, No.5, 883-914.

Logan, G.D. (1995). The Weibull distribution, the power law, and the instance theory of automaticity. Psychological review, Vol. 102, No. 4, 751-756.

Mané, A. & Donchin, E. (1989). The Space Fortress Game. Acta Psychologica, Vol. 71, 17-22.

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

Mazur, J.E. & Hastie, R. (1978). Learning as accumulation: a reexamination of the learning curve. Psychological Bulletin, Vol. 85, 6, 1256-1274.

Myung, I. J., Kim, C. & Pitt, M.A. (2000). Towards an explanation of the power law artifact: Insights from response surface analysis. Memory and Cognition, Vol. 28, 5, 832-840.

Newell, A. & Roosenbloom P. (1981). Mechanisms of skill acquisition and the Law of Practice. In J. R. Anderson (Ed), Cognitive Skills and their acquisition (pp. 1-55) Hillsdale N.J: Erlbaum Associates.

Nosofsky, R.M. & Palmeri, T.J. (1997). An exemplar-based random walk model of speeded classification. Psychological Review, Vol. 104, 266-300.

Palmeri, T.J. (1999). Theories of automaticity and the power law of practice. Journal of Experimental Psychology; Learning, Memory and Cognition, Vol. 25, No.2, 543-551.

Rosenblatt, F. (1961). Principles of neurodynamics. Washington, DC: Spartan.

Schneider, W. (1985). Training high –performances skills, fallacies and guidelines. Human Factors, Vol. 27, 3, 285-300.

Singpurwalla, N.D. & Wong, M.T. (1983). Estimation of the failure rate-A survey of nonparametric methods. Part I: Non-Bayesian methods. Communications in Statistics - Theory and Methods, A12. 559-588.

Snoddy, G.S. (1926). Learning and stability. The Journal of Applied Psychology, Vol. 10. No. 1, 1-36.

StatSoft, Inc. (2001). Electronic Statistics Textbook. Tulsa, OK: StatSoft. WEB: http://www.statsoft.com/textbook/stathome.html.

Restle, F. & Greeno, J.G. (1970). Introduction to mathematical psychology. Reading Mass.: Addison-Wesley.

Uezu, T. & Kabashima, Y. (1996). Asymptotic behaviour of learning from stochastic examples: one-step RSB calculation of the learning curve. Journal of Physics A: Mathematics General. Vol. 29. L55-L60.

Vidulich, M.A., McCoy, A.L. & Crabtree, M.S. (1995). Attentional control and situational awareness in a complex air combat simulation. Paper presented at AGARD Symposium on "Situation Awareness". Published in AGARD report CP-575. Neuilly-sur-Seine, France.

Zandt, T. van (2000). How to fit a response time distribution. Psychonomic Bulletin & Review, Vol. 7. No. 3, 424-465.

This page is intentionally blank

Appendix A Estimating the dead time parameter T_d.

Let $T(i)$, i=1, .., k, be the rank ordering of the k (k n) smallest trial times of a series of n trial times $\{T_1, T_2, \ldots, T_n\}$. We use the minimum trial time $T(1)$ as a first estimate for T_d. Subsequently, we estimate the series of k random errors in the k smallest trial times with:

$$\hat{\varepsilon}_i = T(i) - \hat{T}_d \, . \qquad (26)$$

We approximate the hazard rate $h(\varepsilon_i)$ (i=1, .., k-1) of each error by the empirical hazard rate $\hat{h}_i$ for the i^{th} interval $\Delta\hat{\varepsilon}_i = \hat{\varepsilon}_{i+1} - \hat{\varepsilon}_i$, which can be calculated by (Singpurwalla and Wong, 1983):

$$\hat{h}(\varepsilon_i) \approx \hat{h}_i = \frac{i}{\sum_{j=1}^{i} (n - j + 1) \cdot \Delta\hat{\varepsilon}_j} \, . \qquad (27)$$

By assuming that $\hat{h}_i = \hat{a} \cdot \hat{\varepsilon}_i$ for small i, we make a local linear approximation of the low tail of the empirical hazard function of ε. Using least squares regression, the regression coefficient a is estimated by:

$$\hat{a} = \frac{\sum_{i=1}^{k-1} (\hat{\varepsilon}_i - \overline{\varepsilon}) \cdot (\hat{h}_i - \overline{h})}{\sum_{i=1}^{k-1} (\hat{\varepsilon}_i - \overline{\varepsilon})^2} \, . \qquad (28)$$

Using this linear approximation we can re-estimate the random errors in the k-1 smallest trial times with:

$$\hat{\varepsilon}_i = \frac{\hat{h}_i}{\hat{a}} \, , \qquad (29)$$

and subsequently re-estimate the dead time parameter T_d with:

$$\hat{T}_d = \frac{\sum_{i=1}^{k-1} (T(i) - \hat{\varepsilon}_i)}{k - 1} \, . \qquad (30)$$

We applied the method using k=25, yielding the T_d estimates that are reported in appendix B.

Appendix B Fit statistics for five different learning curve models

Table B.1: Hyperbolic Model $E\left[\boldsymbol{T_n}\right]=\hat{T}_d+\hat{b}\cdot n^{-1}$

Trainee	Parameter Estimates		Fit Statistics			
	T_d [s]	**b [s]**	**MPE [%]**	**MAPE [%]**	**RMSE [s]**	**R^2*100 [%]**
1	4.9	10104.3	92.7	61.6	211.9	48.7
2	4.5	8396.4	56.9	52.0	186.4	67.3
3	3.5	11637.9	101.6	61.8	290.2	59.9
4	4.6	11326.5	97.6	69.1	303.0	20.1
5	3.8	12068.9	103.4	63.3	381.2	60.4
6	6.2	9475.3	65.5	63.3	325.8	6.7
Mean	4.6	10501.6	86.3	61.8	283.1	43.9
SD	0.9	1295.6	18.2	5.1	66.2	22.5

Table B.2: Exponential Model $E\left[\boldsymbol{T}_n\right]=\hat{T}_d+\hat{c}\cdot e^{-\hat{\varphi}\cdot n}$

Trainee	Parameter Estimates			Fit Statistics			
	T_d [s]	**c[s]**	**φ [-]**	**MPE [%]**	**MAPE [%]**	**RMSE [s]**	**R^2*100 [%]**
1	4.9	17.3	-4.5E-04	-14.1	26.5	15.4	17.8
2	4.5	30.6	-9.8E-04	-25.3	36.7	46.1	13.0
3	3.5	32.8	-7.5E-04	-16.4	31.7	29.4	19.3
4	4.6	31.1	-7.2E-04	-16.8	35.5	29.1	16.1
5	3.8	52.0	-1.1E-03	-15.3	33.0	38.0	24.6
6	6.2	64.7	-1.8E-03	-18.9	33.9	61.9	16.6
Mean	4.6	38.1	-9.8E-04	-17.8	32.9	36.7	17.9
SD	0.9	15.6	4.4E-04	3.7	3.3	14.7	3.6

Table B.3: Power Law Model $E\left[\boldsymbol{T}_n\right]=\hat{T}_d+\hat{b}'\cdot n^{-\frac{2}{3}}$

Trainee	Parameter Estimates		Fit Statistics			
	T_d [s]	**b' [s]**	**MPE [%]**	**MAPE [%]**	**RMSE [s]**	**R^2*100 [%]**
1	4.9	942.4	4.2	30.0	19.6	51.1
2	4.5	889.4	-15.3	36.0	29.2	66.3
3	3.5	1242.5	3.4	33.1	26.3	64.4
4	4.6	1235.4	3.8	37.5	37.1	32.7
5	3.8	1511.6	6.4	34.2	39.9	67.0
6	6.2	1218.7	-9.1	38.5	70.1	12.0
Mean	4.6	1173.3	-1.1	34.9	37.0	48.9
SD	0.9	207.8	8.1	2.8	16.2	20.4

Table B.4: Power Law Model $E\left[\boldsymbol{T}_n\right]=\hat{T}_d+\hat{b}''\cdot n^{-0.5}$

Trainee	Parameter Estimates		Fit Statistics			
	T_d [s]	**b'' [s]**	**MPE [%]**	**MAPE [%]**	**RMSE [s]**	**R^2*100 [%]**
1	4.9	287.8	-8.9	25.5	11.8	49.4
2	4.5	289.5	-26.5	35.7	39.5	58.5
3	3.5	406.0	-11.9	29.8	21.1	61.2
4	4.6	408.0	-10.9	33.2	24.6	37.0
5	3.8	535.0	-9.6	30.5	25.5	65.5
6	6.2	437.1	-21.6	36.7	61.5	15.5
Mean	4.6	393.9	-14.9	31.9	30.7	47.8
SD	0.9	85.9	6.7	3.8	16.0	17.2

Table B.5: Progressive average function $E\left[\boldsymbol{T}_{n+1}\right]=\frac{1}{2}\hat{T}_d+\frac{1}{2}\left\langle\boldsymbol{T}\right\rangle$

Trainee	Parameter Estimate	Fit Statistics			
	T_d [s]	**MPE [%]**	**MAPE [%]**	**RMSE [s]**	**R^2*100 [%]**
1	4.9	-5.0	27.2	10.5	60.0
2	4.5	18.5	62.5	28.3	66.3
3	3.5	0.4	35.9	18.6	65.6
4	4.6	-7.1	36.3	26.2	28.9
5	3.8	-3.7	33.8	23.8	68.1
6	6.2	-3.3	50.9	61.0	16.9
Mean	4.6	-0.1	41.1	28.1	51.0
SD	0.9	8.6	11.9	15.8	20.3