Non-random Deviations from the Power Law of Practice

Hiroaki Suzuki [1] (susan@ri.aoyama.ac.jp), Hitoshi Ohnishi [2] (ohnishi@nime.ac.jp),
[1] Department of Education, Aoyama Gakuin University, Japan
[2] National Institute of Multimedia Education, Japan

Abstract

Many studies on the acquisition of simple skills have claimed that improvements in performance follow the power law of practice. However, it is also well-known that during long-term practice there are fluctuations such as plateaus, slumps, and spurts. We analyzed a relatively long-term learning process in a simple assembly task to objectively examine fluctuations during the learning process. We first applied time-series analysis based on the state space method to the task-completion time. The analysis revealed that the power law of practice is only a first approximation, and that fluctuations in the learning curve reflected the long-term trends. Second, we focused on one of the fluctuations, and carried out microscopic analysis to find what generated the slumps and breakthroughs. We found that, contrary to dominant skill-acquisition models such as ACT-R (Anderson, 1993), the slump was attributable to the mismatch between the level of skill and the external environment that the skill operates. This analysis suggests that, to fully elucidate the processes and mechanisms of skill acquisition, attention should be paid not only to the internal mechanisms, but also to the external environment that the skills operate.

Introduction

We improve skills by practice. The more we practice an action, the less time it takes to carry it out. Our skills improve greatly at the beginning, but the rate of improvement gradually slows down. We cannot speed up action at a constant rate. This phenomenon has attracted the interest of many researchers. Some have focused on simple physical skills such as card sorting and cigar rolling, while others have focused on complex cognitive skills, such as addition, and geometry proof (Crossman, 1959; Neves & Anderson, 1981).

Most researchers in this field generally agree that learning curves follow the power law of practice. This law is expressed as:

$$T = NP^c,$$

where T is the time to perform the task, P is the amount of practice, and N is the time to perform the first trial. c is a negative number representing the learning rate. This law states that the time to complete a task decreases in proportion to the power of the number of trials.

If we draw a figure whose abscissa is the logarithm of the number of practice and the ordinate is the logarithm of the time to complete the task, we can obtain a straight descending line. This log-log plotting nicely fits the layman's observation as well as data obtained from scientific research. First, the execution time becomes faster with practice. Second, the rate of speed-up gradually (logarithmically) decreases. We can make marked progress in the initial stage, but such progress can never be obtained in the middle or later stages.

Many theorists have tried to develop models of this phenomenon. ACT-R proposes that the large amount of speed-up in the early stage is the result of knowledge compilation, and that the later improvement is due to knowledge tuning (Anderson & Lebiere, 1998). The initial knowledge enabling an action in ACT-R takes the declarative form. When doing an action, this declarative knowledge must be interpreted by domain-general productions. After several practices, however, the knowledge compilation mechanism transforms the declarative knowledge to procedural. A production rule, i.e., a condition-action pair that is directly executed by the system, is created during this process. Thus, knowledge compilation dispenses with the laborious interpretation required at the initial stage, which largely reduces the time to execute an action. Further practice tunes production rules by adjusting the strengths of productions on the basis of the success probability, the value of achieving the goal, and the cost of performing the production.

SOAR provides a similar account of the learning curve. SOAR's chunking mechanism converts multiple cognitive steps into a single production rule. Since small useful chunks are likely to be created at the early stage, the time for the task greatly reduces. However, the probabilities of new chunks being created gradually decrease as learning continues, which leads to the slow speed-up rate.

Although ACT-R and SOAR provide coherent computational accounts of the power law, this kind of macroscopic characterization conceals well-known facts in learning, i.e., plateaus, slumps, and breakthroughs. Anyone who has experienced extensive practice can remember periods where no performance gains were obtained even after repeated practice. In some cases, performance may deteriorate to a certain extent. Some people have fortunately had experiences with breaking a seemingly lasting slump.

Kimura's study (1999) provides valuable information about this. Kimura analyzed learning curves obtained from more than 10,000 times practices in origami (Japanese paper folding). Subjects greatly reduced the time required to complete the origami from several minutes at the beginning to less than half a minute at the end. He found that there were far more valleys and hills than theoretically estimated. Furthermore, a detailed analysis revealed two important regularities. First, breakthroughs follow relatively long slump periods. Second, after outstanding records are achieved, performance often deteriorates. Kimura's analysis as well as that by Seibel (1963) suggest that fluctuations do not occur randomly, but that fluctuations may be manifestations of some important changes occurring in learning processes.

The second problem in the previous studies is that researchers have not paid enough attention to the environment where the actions are carried out. As far as motor skills are concerned, they are not carried out mentally, but physically. This means that skills presuppose a specific environmental setting. In other words, skills are joint products of both internal mechanisms and the environment. Therefore, it is not sufficient to only analyze internal mechanisms. We should also analyze skills with respect to the environment where the skills are demonstrated.

Thelen and Smith (1994) nicely showed this intertwined nature of action. Whereas six-month-olds do not usually exhibit the stepping reflex, they can perform alternate stepping in the bathtub with water at waist level. This means that some internal mechanisms presuppose specific environmental settings. If an appropriate setting is not available, potentially executable mechanisms cannot work well.

Since skill learning and motor development share many features in common, we tentatively formulated a hypothesis about slumps and breakthroughs. Slumps are caused by mismatches between developing internal mechanisms and environmental settings. People refine and strengthen their internal mechanisms by practice. However, environments that have supported less developed skills may sometimes become inappropriate for more advanced skills, which results in slumps. During further practice, people occasionally find appropriate settings that then support the advanced skills they are developing, which leads them to breakthroughs.

The purpose of the present study was two-fold. First, we examined whether learning curves involved components that the power law of practice did not explain. To do this, we applied the state space model to the time-series data obtained from a simple assembly task.

Second, to identify sources of slumps, plateaus, and breakthroughs, we carried out microscopic analysis on a part of learning where a slump and deterioration were followed by a breakthrough. In this analysis, we focused on the interrelation between internals and externals. We tried to find whether mismatches were the main sources of slumps, and whether it was crucial to reconfigure the environment to achieve breakthroughs.

Overview of experiment

Method

Subject: A female undergraduate student participated in the study. She was paid 1000 yen per day. In addition, to encourage improvement, we told her that she would receive a bonus if her best time spent on tasks on any day was 10% less than that spent on a previous day.

Procedure: We provided her with a Lego block model. The model roughly had three parts, i.e., a body and left and right wings. The body consisted of three medium-sized blocks, and each wing of short and long blocks. All the blocks were of different colors.

We placed six short and medium, and three long blocks in front of the subject. We told her to replicate the model by assembling the blocks in front of her. The experimenter told her to start and measured the time taken in each trial. Fifteen trials constituted a session. After one session had finished, the subject was given a one-minute break. This cycle lasted until about an hour had passed. She continued this for twelve consecutive days.

Results

The total number of trials amounted to 2325 (155 sessions) in 12 days. The task completion time was 38.6 sec in the initial trial, and decreased to 2.83 sec in the 2303rd trial on the last day. The number of trials also indicated the subject's improvement. While she performed 135 trials within an hour on the first day, she was able to perform 210 trials on her last day.

The decrease in the task-completion time was remarkable on the first day. It reduced from 38.6 sec to 12.59 sec within the first session. The best time on the first day was 5.98 sec. Thus, the subject reduced the task-completion time by about 32 seconds on the very first day. The rate of decrease becomes more and more moderate during practice. The decrease was about 1 sec on the second day, about 0.5 sec on the third and fourth days, and 0.1 to 0.3 sec on the remaining days.

Since this pattern of change seemed to fit the power law, we first employed the power law model. The relation obtained between the task-completion time and the number of trials is shown in Eq. (1) and Fig. 1.

$$y_n = 28.08n^{-0.269}, \tag{1}$$

where y_n stands for the observed task-completion time for the n-th trial.

Although the model may appear to fit the data, we carried out more careful analysis to find systematic deviations from the estimates by the power law of practice. We considered a sequence of trials to be a plateau if the best time for any given trial could not be improved over more than 50 trials. We found 12 plateaus, whose length ranged from 58 to 606 trials (the average was 164.2). This demonstrates that the power law of practice is only a first approximation, and the learning curve involves many plateaus.

Analysis with state space model

Typical approaches to estimating the trend, which represents long-term or macro movement in a time series, have employed parametric-polynomial-regression or generalized-linear-regression models. The quality of the estimates have depended on the appropriateness of the assumed model. If this is inappropriate, the model cannot be fitted to the data or used to extract the essential structure of the data. The power law model, a kind of generalized linear model, assumes that the task-completion time will decrease monotonically with the number of trials. Therefore, the increases in the task-completion time are regarded as random errors.

However, this observation suggests that fluctuations around the power law line are not fully random. Thus, it is necessary to employ a model that treats fluctuations not as random errors, but as essential components of time-series data. It is also advantageous for the model not to require strong assumptions about its class.

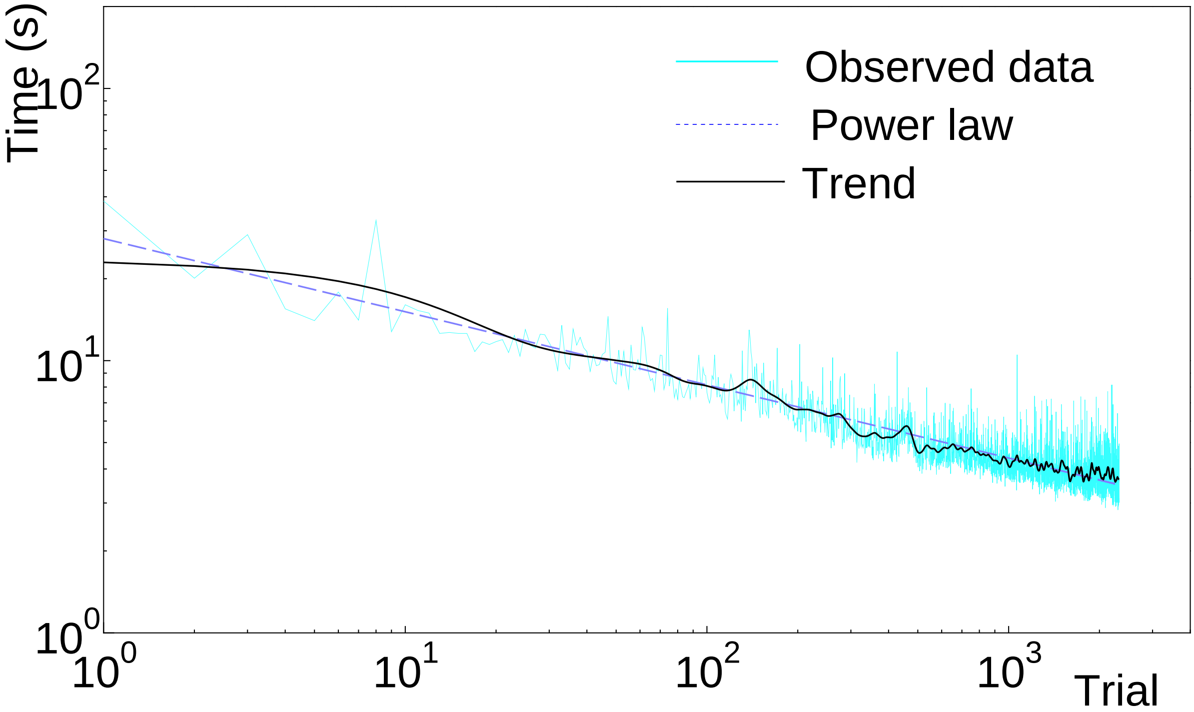

Figure 1: Observed task-completion time, fitted power law model, and estimated trend

Of the many models available, we applied the state space model (Kitagawa & Gersch, 1984). This model can capture non-random fluctuations as trends of data with weaker assumptions. We assumed the task-completion time of the n-th trial y_n would be decomposed into a trend component, t_n, and an irregular component, w_n, as

$$y_n = t_n + w_n, \tag{2}$$

where $w_n \sim N(0, \sigma^2)$. Trend component t_n was assumed to be a smooth stochastic process and satisfies Eq. (3).

$$t_n = 2t_{n-1} - t_{n-2} + v_n, \tag{3}$$

where $v_n \sim N(0, \tau^2)$. This assumption is weaker than the assumptions for polynomial and generalized linear models. Therefore, the state space model can be used to estimate the trend flexibly.

These models can be represented as a state space model given by Eqs. (4) and (5).

$$\begin{aligned} x_n &= F x_{n-1} + G v_n & (4)\\ y_n &= H x_n + w_n, & (5)\end{aligned}$$

where

$$x_n = [\ t_n \quad t_{n-1}\ ]^t,$$

$$F = \begin{bmatrix} 2 & -1 \\ 1 & 0 \end{bmatrix}, G = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, H = [\ 1 \quad 0\ ],$$

$$Q_n = \tau^2, R_n = \sigma^2.$$

The state, x_n, can be estimated with the following Kalman filter and smoothing algorithm if parameters σ^2, τ^2 are given. The Kalman filter is used to estimate unobservable x_n by prediction based on the dynamics of the system (Eq. (4)) and by filtering (adjustment) based on observation (Eq. (5)).

The specific computation for prediction is given by Eq. (6) and that for filtering is given by Eq. (7)

$$\begin{aligned} x_{n|n-1} &= F x_{n-1|n-1} \\ V_{n|n-1} &= F V_{n-1|n-1} F^t + G Q_n G^t. \end{aligned} \tag{6}$$

$$\begin{aligned} K_n &= V_{n|n-1} H_n^t (H_n V_{n|n-1} H_n^t + R_n)^{-1} \\ x_{n|n} &= x_{n|n-1} + K_n (y_n - H F x_{n|n-1}) \\ V_{n|n} &= (I - K_n H_n) F V_{n|n-1}. \end{aligned} \tag{7}$$

Using the outputs of the Kalman filter, smoothed state x_n is estimated by

$$\begin{aligned} A_n &= V_{n|n} F_{n+1}^t V_{n+1|n}^{-1} \\ x_n &= x_{n|N} = x_{n|n} + A_n (x_{n+1|N} - x_{n+1|n}) \\ V_{n|N} &= V_{n|n} + A_n (V_{n+1|N} - V_{n+1|n}) A_n^t \end{aligned} \tag{8}$$

Parameters σ^2, τ^2 were estimated by the maximum likelihood method. The log-likelihood function of the model is given as

$$\begin{aligned} \ell(\theta) = -\frac{1}{2} \{ & N \log 2\pi + \sum_{n=1}^{N} \log r_n \\ & + \sum_{n=1}^{N} \frac{(y_n - H_n x_{n|n-1})^2}{2r_n} \}, \end{aligned} \tag{9}$$

where $r_n = H_n V_{n|n-1} H_n^t + R_n$. The parameters are obtained by maximizing Eq. (9) with respect to these parameters.

The estimated trend component is plotted in Fig. 1. If the learning process follows the power law of practice, the trend component becomes identical to the power law line. However, the trend component fluctuates around the power law line.

In contrast to the power law model, the state space model can capture non-random fluctuations observed in the time-series data. For example, there is a long plateau between 300 to 600 trials. In the 314th trial, the subject achieved the best time but this was not broken until about the 500th trial. Figure 2 plots the observed data and estimates produced by the state space model and the power law model. As is obvious from the figure, estimates obtained by the power law model monotonically decrease, while the state space model nicely follows the observed data. In addition, the state space model is not affected by one-shot flukes or bad luck.

These analysis revealed that, contrary to the prediction by the power law model, the learning process involves a number of plateaus and breakthroughs that follow. The state space model can a give better description of the observed learning curve.

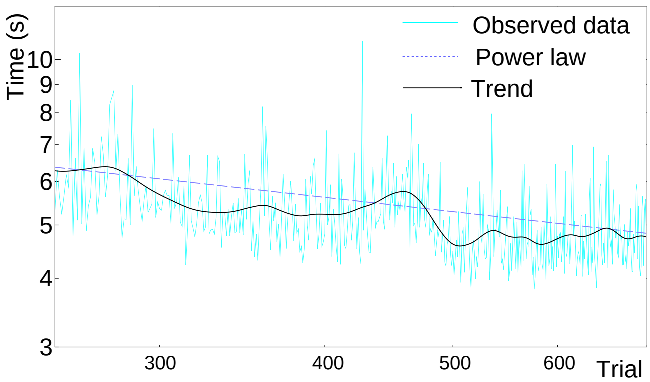

Figure 2: Observed task-completion time, fitted power law model, and estimated trend at interval between 300-th and 600-th trials.

Microscopic cognitive analysis

We focused on a slump, a regression, and a breakthrough that appeared from the middle of the third day to the fourth. Before proceeding to a detailed analysis, it might be of some help to describe the general pattern in the subject's assembly

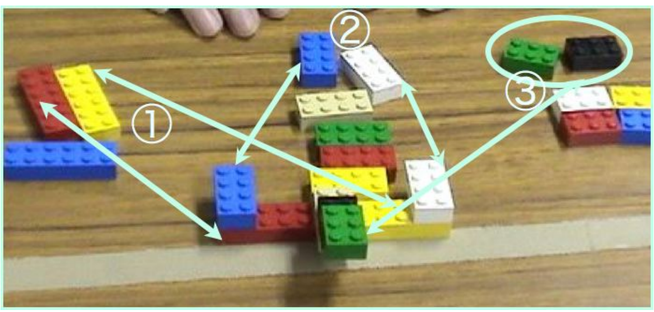

Figure 3: Pattern for subject ' s assembly of Lego.

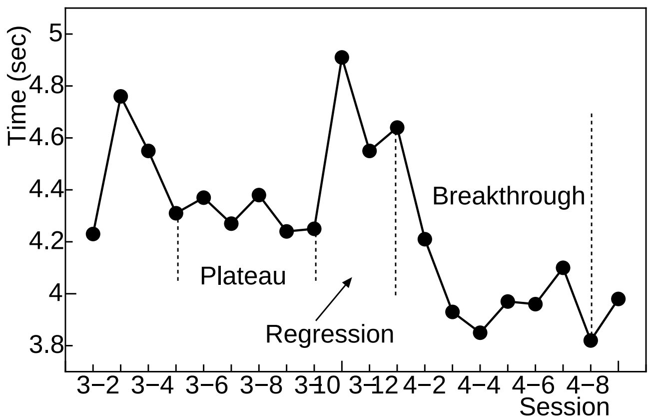

Figure 4: Slump, regression, and breakthrough observed on third to fourth days.

of the blocks. As shown in Fig. 3, she first took and aligned two wing blocks. Second, she connected the third block on the top of the right wing, then the fourth block on the top of the left wing. Finally, she constructed the head part by consecutively connecting the three blocks, and attaching this to the center of the wing part. She did the task in this way from the very beginning to the end of the entire training sessions.

The best times for the sessions were 4.3 to 4.4 sec in the middle of the third day . This lasted for about 90 trials. However, in session 10 on the third day, the best time suddenly increased to 4.8 seconds. This lasted for three sessions (45 trials). However, after this period of regression, the assembly time suddenly decreased to less than 4 sec. As we can see from Fig. 4, there are three distinct periods, the plateau (from 3-5 to 3-10), regression (from 3-11 to 4-1), and breakthrough (from 4-2 to 4-6).

We measured the time it took to connect each block to identify where time was being reduced the most. This analysis revealed that the most contributing part to the reduction of the total assembly time was the one where the fourth block was connected to the top of the left wing. During this period, the time to connect this block decreased by 0.3 seconds. In addition, the fluctuation pattern for time in this part approximately paralleled that of the total assembly time.

We did microscopic analysis on what happened when these blocks were connected. Before the fourth block was connected to the left wing, the subject held the left wing block with her left middle finger. Then, she took the fourth block with her right hand and connected it to the top of the left wing. This caused trouble because her ring and little fingers were at the very point where the fourth block was to be connected. To avoid these difficulties, the subject employed two strategies. The first was to bend her fingers, and the second was to open them outward (See Fig. 5). Both strategies produced sufficient space to connect the fourth block to the top of the left wing.

The subject had mainly used the finger-bending strategy from the beginning of practice. This strategy produced 80% of the best times for the sessions, although she occasionally used the finger-opening strategy as well. However, during the regression period, she used the finger-opening strategy more

Figure 5: Finger-opening vs. finger-bending strategies.

Table 1: Proportions of finger-opening strategy (FOS) in sessions and percentages of best times produced by finger-opening strategy

	% of FOS	Min by FOS (%)
slump	11.4	26.7
regression	31.1	50
breakthrough	65.0	100

frequently. She used this strategy more than 30% of the trials. In addition, it produced 50% of the best times in the sessions in this period. Furthermore, in the breakthrough period, she used the finger-opening strategy for almost two thirds of the trials. All the fastest times in the sessions in this period were produced when she used this strategy.

This microscopic analysis suggested that fluctuations in the learning process can be accounted for in terms of strategy competition. In the initial plateau, the finger-bending strategy suited the subject's level more for some reason. However, her occasional use of the finger-opening strategy at the same time produced the fastest time in several sessions. This might have gradually led her to shift her strategy choice. Consequently, the two strategies became competing. Generally, when more than one strategy is simultaneously activated, the process is disrupted. In this task, the subject had to choose one of the two and adapt her actions before and after the strategy was executed. This would be the main reason the assembly time increased in the regression period. However, by being frequently used, the finger-opening strategy would become more sophisticated and automated. This was evidenced by the fact that in the breakthrough period she mainly relied on and produced all the fastest times using this strategy.

It is important to note that neither strategy involved the main action of connecting the fourth piece to the top of the left wing. Rather, these strategies dealt with preparing an environment for the fourth block to be connected to the top of the left wing. The subject's little finger as well as her hand constituted an environment where the main action was carried out.

If the environment is suitably arranged, the main action should be able to be carried out smoothly. Otherwise, it will take longer to execute it.

Discussion

We aimed at examining whether deviations from the power law of practice are random and at revealing the interrelationship between skills and the environment in the present study.

A detailed analysis of the time series for the task-completion time revealed that there were non-random deviations from the power law. There were a dozen plateaus and breakthroughs that followed. We then applied the state space model to the time-series data. The analysis revealed that there was a systematic trend component. A model with the trend produced better estimates of the subject's performance.

We next carried out microscopic cognitive analysis on a process involving a plateau, regression, and breakthrough. We found competition between the two strategies and shift from one to the other during this process. The pattern of shift was nicely correlated with the fluctuations in the task-completion time. It is important to note that these strategies worked in the background, so that the subject could rapidly carry out her main action. Therefore, a shift in strategy should be considered as a reconfiguration of the environmental setting where the main action is performed.

Our analysis also revealed that slumps and regressions are sometimes caused by mismatches between main actions and the environment. When one wants to quickly connect a block during an assembly task, what one has to do is not only to bring it quickly to the connection point, but also to prepare the setting where the action can be carried out swiftly. The preparation involves creating enough space to connect the block. For our subject and her level of expertise at least, it was easier for her to open her fingers outward than to bend them. This is why she shifted to the finger-opening strategy.

However, the shift required many things to be done. First, one has to make the new strategy more sophisticated. Second, it is also necessary to coordinate the strategy with the actions carried out before and after. Finally, one has to inhibit a highly automated old strategy for the new one to work. These are not easy tasks and they take time. This is why the strategy shift was gradual rather than abrupt.

The lesson from this analysis is that to fully elucidate the process and mechanisms of expertise, one should focus more on the reconfiguration of the environment. Previous models such as ACT-R (Anderson, 1993), SOAR (Newell and Rosenbloom, 1981), and component theory (Speelman and Krisner, 2005) have mainly dealt with actions that are directly concerned with achieving tasks. But, since skills are joint products of both internal mechanisms and the environment, it is important to synthesize the analyses of internals and externals.

References

Anderson, J. R. *Rules of the Mind*. Hillsdale, NJ: LEA.

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

Kimura, I. (1999) Plateaus, slumps, and the power law of practice: A paper-folding study with initial results. In Proceedings of the Second International Conference of Cognitive Science Society, 394 – 399.

Kitagawa, G. and Gersch, W. (1984). A smoothness priors-state space approach to the modeling of time series with trend and seasonality. Journal of the American Statistical Association, Vol. 79, No. 386, pp. 378–389.

Neves, D. M. and Anderson, J. R. (1981) Knowledge compilation: Mechanisms for automatization of cogntive skills. In J. R. Anderson (Ed.) *Cognitive skills and their acquisition*. Hillsdale, NJ: LEA.

Newell, A. and Rosenbloom, P. S. (1981) Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.) *Cognitive skills and their acquisition*. Hillsdale, NJ: LEA.

Seibel, R. (1963) Discrimination reaction time for a 1,023-altenative task. *Journal of Experimental Psychology*, 66, 215 – 226.

Speelman, C. P. and Kirsner, K. (2005) Beyond the Learning Curve: The Construction of Mind. Oxford, UK: Oxford University Press.

Thelen, E. and Smith, L. B. (1994) *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT.