Evidence of an optimal error rate for motor skill learning

Naser Al-Fawakhiri[1], Sarosh Kayani[1], & Samuel D. McDougle[1,2]

1. Department of Psychology, Yale University, New Haven CT 06520
2. Wu Tsai Institute, Yale University, New Haven CT 06520

Correspondence: Naser Al-Fawakhiri
naser.al-fawakhiri@yale.edu
Department of Psychology
Yale University
100 College Street, New Haven, CT, 06510

Abstract

When acquiring a motor skill, learners must practice the skill at a difficulty that is challenging but still manageable in order to gradually improve their performance. In other words, during training the learner must experience success as well as failure. Does there exist an optimal proportion of successes and failures to promote the fastest improvements in skill? Here, we build on a recent theoretical framework for optimal machine learning, extending it to the learning of motor skills. We then designed a custom task whose difficulty dynamically changed along with subjects' performance, constraining the error rate during training. In a large behavioral dataset, we observe evidence that learning is greatest at around a ~30% error rate, matching our theoretical predictions.

Author Summary

Practicing a motor skill involves successfully performing intended movements and learning from mistakes. Is there an optimal proportion of mistakes during training that leads to fast and efficient skill learning? Inspired by recent theoretical work on principles of machine learning, we mathematically derive an "optimal error rate" for a simple motor skill and then experimentally validate our predictions. We find both theoretical and empirical evidence suggesting that ~30% is the optimal error rate for motor learning, which has practical implications both for rehabilitation and for sports coaching and training.

Introduction

Acquiring a new motor skill involves reducing errors [1]-[3]. Consider learning to shoot free throws in basketball: After a miss, the player uses the observed error information (e.g., "I missed to the right") to fine tune future attempts (e.g., re-orienting the shot slightly left). Such performance errors during motor skill learning drive improvements in both accuracy and precision, and given enough time and practice can lead to automaticity and expertise [4]-[6].

Importantly, one's error rate during training can relate to if, and how well, they learn a skill. In the extreme, if one *never* experienced success, their motor system would have no template upon which to improve their skill [7]. On the other hand, if one *always* experienced success (e.g., the hoop moved to wherever the ball was thrown), learning would be superfluous. Thus, it is reasonable to assume that there is a "sweet spot" error rate that may generate the most efficient learning process [8]. Indeed, in the context of simple non-motor skills (e.g., perceptual decision-making) it has recently been proposed that an optimal proportion of success and failure – an optimal "error rate" – can maximize the learning rate of a perceptual skill, such as

binary categorization of sensory input [9]. Here we apply this logic to the domain of sensorimotor skill learning.

In any task, errors become more likely as the difficulty of the task increases. Since the difficulty of the task can often be externally modulated, controlling an individual's error rate during training requires dynamically controlling the difficulty of the task in response to their performance, exemplified by classic psychophysical methods like "staircasing." But how does task difficulty influence learning? In an elegant framework recently proposed by Wilson and colleagues [9], it was shown that learning using a gradient descent algorithm mathematically specifies an optimal error rate with very few assumptions needed. In this recent influential study, the authors considered the context of a binary discrimination task and, using their framework, predicted an optimal error rate of ~15%, a value which was confirmed by multiple model simulations over multiple binary decision-making tasks.

Can a similar framework be applied to motor skill learning? The first relevant issue is the relationship between motor skill learning and gradient descent. It has been suggested that motor skill learning, at least in the form of the gradual improvements in skill that reduce errors, does indeed operate via a gradient descent process [7,10]. For instance, Pierella et al. [10] trained human participants in a skill that required them to use a novel mapping of upper body movements to control a visual cursor. As they learned the appropriate internal model, subjects' behavior for this task closely matched simulated data generated by a gradient descent process, suggesting that the formation of internal models during motor learning can, in some cases, be well approximated as a form of gradient descent. State-space models, a common modeling framework for motor learning, employs a delta rule to minimize sensory prediction errors [11,12], thus employing a form of gradient descent. This minimization of sensory prediction

errors during motor learning is thought to be subserved by the cerebellum, which implements neural algorithms that gradually reduce error signals [13, 14].

Here we hypothesized that the elegant framework proposed by Wilson and colleagues [9] might apply to motor skill learning, even though motor errors are both signed and continuous (rather than binary and unsigned). First, we analytically extended their model to demonstrate that a roughly 30% error rate may be the optimal error rate for motor skill learning, given some basic assumptions. After specifying this prediction, we developed a customized "Pong" task that required precise online visuomotor control and collected behavioral data on the task from a large sample of human participants (N = 194). Using a custom dynamic error-control algorithm, we modified the size of people's Pong paddle on a trial-by-trial basis to impose different error rates on their performance and measured their improvements in the task. Strikingly, the results of the behavioral experiment appeared to support our predictions – we found reliable evidence that individuals with an error rate around ~32% showed the most learning in the task, in line with our model simulations. These findings provide a theoretically grounded specification of a putative optimal error rate for motor learning, and may thus inform both practical applications (e.g., sport pedagogy, clinical rehabilitation protocols) and theories concerning the underlying neural algorithms for motor skill learning.

Results

A mathematical derivation of the optimal error rate for motor skill learning

The derivation below largely parallels the derivation of the optimal error rate for binary discrimination tasks, as outlined by Wilson and colleagues [9]. In an idealized motor learning task, an agent prepares an action $\boldsymbol{h}$ in response to some set of internal conditions $\boldsymbol{p}$ and external stimuli $\boldsymbol{x}$. Thus,

$$h = \Phi(\vec{p}, \vec{x}, \phi) \quad (1)$$

Where $\Phi(\cdot)$ represents the function which transforms ***p*** and ***x*** into motor actions, and ϕ is a set of internal, tunable free parameters trained by learning.

Whenever an action is executed, it is evaluated in terms of the discrepancy between the actual movement outcome ***h*** and the intended outcome ***i***. This discrepancy is a motor error ***e***

$$e = h - i \quad (2)$$

Here we restrict our attention to one dimensional errors (i.e., the vertical axis of a pong workspace), although this process could in theory be generalized to higher dimensions. We further assume that errors are normally distributed around the target action and that both positive and negative signed errors are equally likely. Thus, ***e*** can be rewritten as a Gaussian with zero mean and a standard deviation of σ

$$e \sim N(0, \sigma) \quad (3)$$

Both laboratory and naturalistic motor tasks typically have some degree of error tolerance, Δ, in order to still permit task success given noise. This tolerance can take many forms: the radius of the hole in golf, the radius of the target in a center-out reaching task, or, in our Pong task, the center-to-edge length of the paddle. Under these conditions, the error rate at any point in the task can be formalized as:

$$ER = 2\int_{-\infty}^{\Delta} \mathrm{p}(\mathrm{e}|0, \sigma)\ \mathrm{de} = 2\mathrm{F}(-\Delta/\sigma) = 2\mathrm{F}(-\beta\Delta) \quad (4)$$

where *p*(·) is the standardized distribution, here a normal distribution, $F(x)$ is the cumulative density function of this probability distribution, and $\beta = 1/\sigma$ is the agent's motor precision. In our case, β would be the inverse of the root-mean-squared-error (RMSE) centered around the intended movement ***i***.

Since ***e*** is directly related to ***h***, the goal of learning is to tune the parameters ϕ responsible for generating ***h*** so that errors ***e*** are minimized and ***h*** is as close to ***i*** as possible. We here assume that improvements in Pong performance abide by gradient descent to tune these parameters [10]. We can thus formalize the process of learning over time *t* as:

$$\frac{d\phi}{dt} = -\eta \nabla_{\phi} ER \tag{5}$$

where η is the stepwise learning rate of the gradient descent algorithm. Given Equation 4, the gradient of the error rate can be rewritten via the chain rule:

$$\nabla_{\phi} ER = \frac{\partial ER}{\partial \beta} \cdot \nabla_{\phi} \tag{6}$$

The first term of the right-hand side is the only term which depends on Δ, the task success tolerance (which can be experimentally controlled), as the second term only reflects how precision changes with the tuning of the relevant parameters. Thus, to maximize the learning rate in Equation 5, we must maximize $\frac{\partial ER}{\partial \beta}$ by optimizing the task success tolerance Δ at each time point *t* in training (see *Methods*). The tolerance Δ* which maximizes the learning rate $\frac{\partial ER}{\partial \beta}$ can be plugged into Equation 4 to yield the optimal error rate $ER^{\star}$. Given that motor errors are normally distributed in most tasks, the optimal error rate $ER^{\star}$ can be computed:

$$ER^{\star} = 1 - erf\left(\frac{1}{\sqrt{2}}\right) \approx 0.3173 \tag{7}$$

In other words, in this framework, the optimal error rate for motor learning is 31.73%, and an optimal training schedule should aim to maintain performance at a 68.27% success rate.

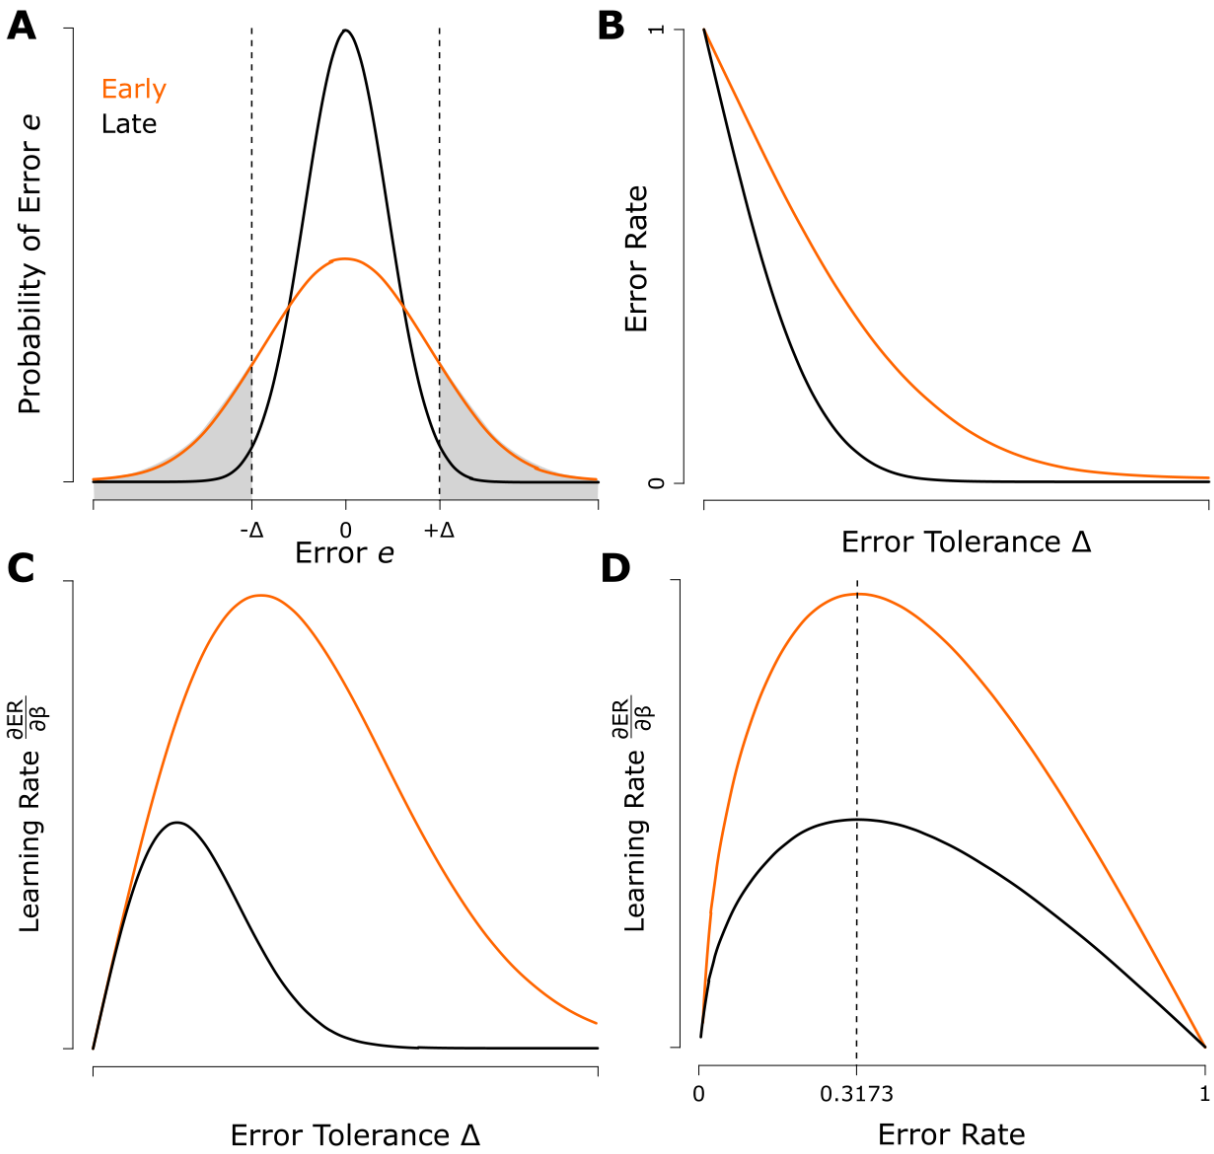

Figure 1. Derivation of an optimal error rate for motor learning. Figure largely adapted from Wilson et al. [9]. (A) Probability distributions of a given motor error *e*, within a certain error tolerance Δ. Shaded area represents the probability of a failure, or the error rate. Across training, the magnitude of errors shrinks (black line), resulting in increased precision β. (B) As the error tolerance Δ increases (i.e. decreasing difficulty), fewer actions result in failures, leading to a decrease in the error rate. (C) The learning rate is given by the derivative of the error rate (ER) with respect to the agent's precision β. Maximal learning occurs at an error tolerance Δ which changes with skill/training (β). (D) Same as C but replotted with the corresponding error rates on the x-axis. The optimal error rate remains constant, regardless of precision β, at ~31.73%.

Skill improvement in a modified Pong task is optimized when performance is maintained near a 30% error rate

We aimed to test our theoretically predicted “optimal error rate” for motor learning using a motor skill task fashioned after the video game Pong. In this task, participants (N=194) controlled the vertical position of a rectangular paddle using a mouse or touchpad in order to intercept a small white puck that moved across the screen at various angles. If participants successfully intercepted the puck with the paddle, their movement was considered successful and they were awarded 10 points. If the puck intersected the vertical axis of the paddle without striking the paddle, a red “0” flashed on the screen indicating failure.

In order to scale the difficulty of the task with participants’ evolving skill level, and thus maintain as close to a fixed error rate as possible during training, the paddle’s size dynamically changed across trials according to the distribution of participants’ errors from the center of the paddle over the most recent 20 trials (see *Methods* for details). Participants were randomly sorted into 4 groups, each with an imposed error rate: 15%, 30%, 50%, and 65%.

After an initial 20-trial calibration phase in which the puck was calibrated to move at a reasonable speed (see *Methods*), and 10 baseline trials where both the puck speed and paddle size were held constant, participants completed 70 training trials in which the paddle size varied with their performance. To track performance improvements across the task, we computed the root-mean-squared-error (RMSE) of the puck’s final position relative to the center of the final position of the user-controlled paddle for every 10-trial bin. As a proxy for precision, we inverted this RMSE to yield β; β is low when precision is low, and high when precision is high. Since every participant entered the task at various degrees of skill, we corrected these β using the 10 trial baseline in order to measure each participant’s *improvement* (i.e., learning) across the length of the task. A “learning index” was quantified as the average β over all 70 training trials.

Baseline rates of successful puck interception were not different between groups and the average success rate in the baseline was 77% (SD:17%). The dynamic updates to the paddle size (see *Methods*) successfully constrained participants' error rates across groups (Fig. 2). However, we note that the 65% error rate group did have modestly but significantly greater success than the intended 35% success rate (t(48)=3.03, p=0.004), averaging 41.6% success (SD: 6.1%) after the baseline. Participants' true error rates, defined as the percentage of miss trials during the 70 trial baseline phase, on average undershot the intended error rate by 3.1% (SD: 6.6%), or 2.2 (SD: 4.6) trials.

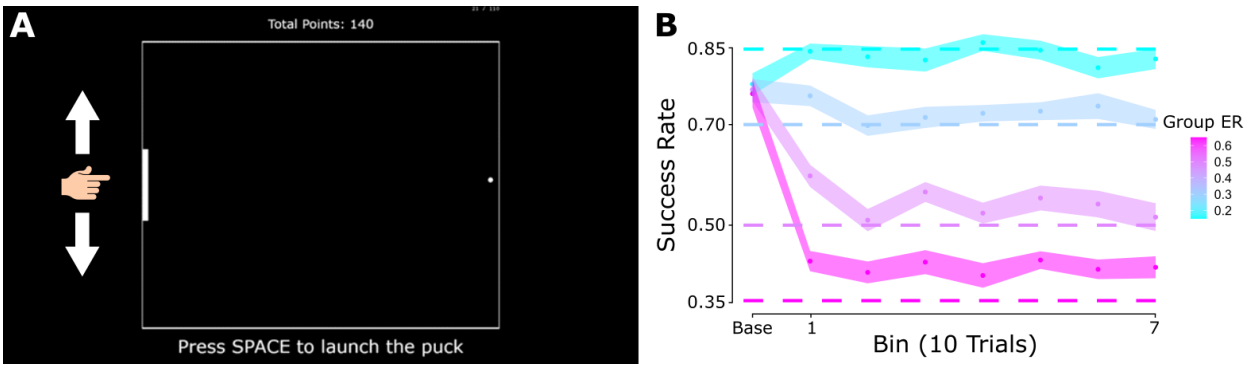

Figure 2. Task description. (A) Participants played a modified version of Pong, where they could control the paddle's vertical position with their mouse in order to intercept the moving puck. (B) The paddle's length was dynamically controlled on each trial by an algorithm designed to control the error rate across trials. Each participant was randomly assigned to one of four groups, each with a specific error rate ("ER"; 15%, 30%, 50%, 65%). The algorithm was largely successful in limiting the success rates of participants in each group. Error shading reflects 1 S.E.M.

Regardless of error rate group, participants improved their precision (inverse RMSE, β) during the training phase (Fig. 3A). Participants' precision in the final 10 trials was significantly greater than their precision in the first 10 training trials (t(386)=2.36, p=0.02, d=0.24) and participants overall improvement, or Learning Index (LI; β averaged across all 70 training trials) was significantly greater than zero (t(193)=4.33, p=2.4x10^{-5}, d=0.31).

Different error rate groups exhibited different time courses and degrees of improvement (Fig. 3B, C). As predicted, the 30% error rate group exhibited the greatest LI over all four groups: Their LIs were significantly greater than the 15% group (t(89)= 3.11, p=0.003, d=0.65), and marginally greater than the 65% group (t(87)=1.58, p=0.12, d=0.33), but not significantly greater than the 50% group (t(94)=0.98, p=0.32, d=0.20).

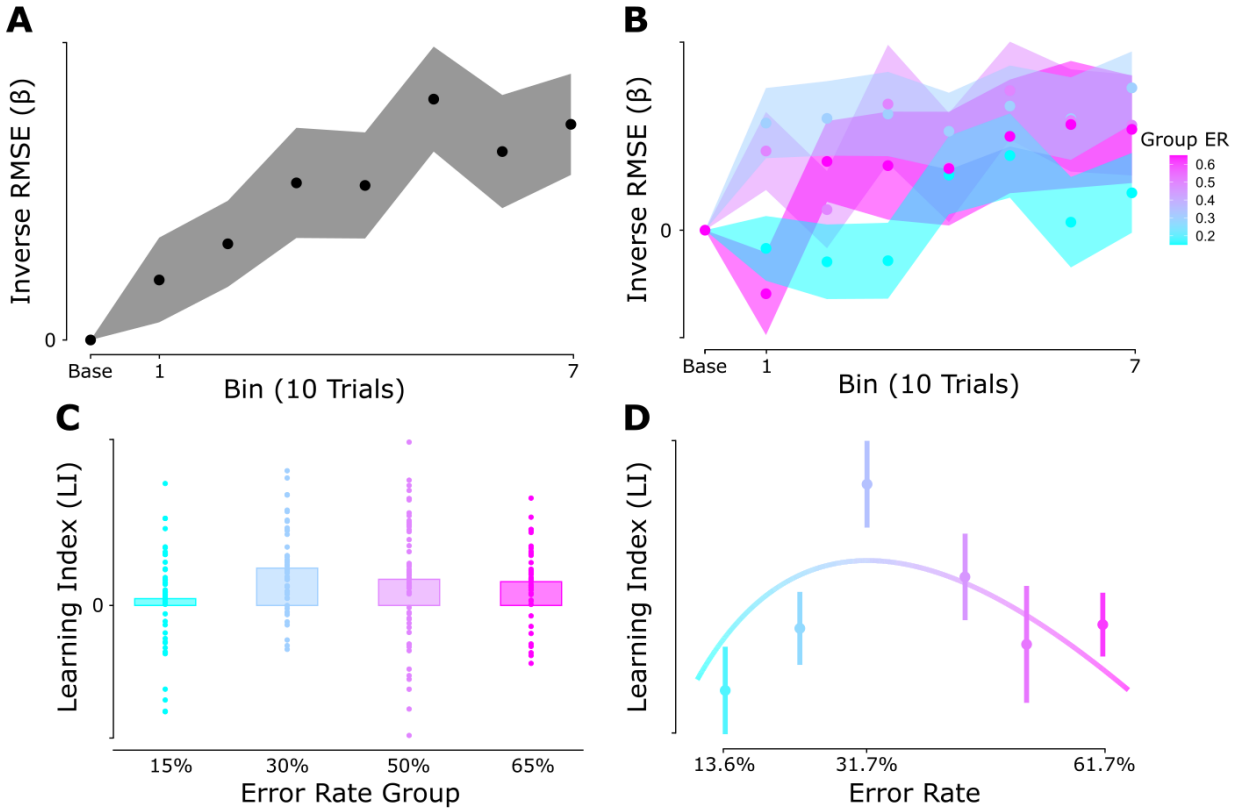

Figure 3. Optimal learning when training has an error rate of ~31.7%. (A) All participants significantly improved over the course of training. β values are the baseline-corrected inverse RMSE for each 10-trial bin (arbitrary units). (B) Each group displayed different degrees and time courses of improvement. (C) Learning index reflects the average β across all 7 training bins for each subject. The group with the intended learning rate of 30% exhibited the greatest learning. (D) Subjects were binned into sextiles based upon their true error rates during training. The sextile with error rate 31.7% had the highest LI, corresponding to the theoretical optimal error rate. The smooth curve reflects the theoretically predicted learning rates for each error rate. Shading in panels A & B and error bars in panels C & D reflect standard errors of the mean (SEM).

However, participants' group assignments do not reflect their true error rates, as our algorithm cannot perfectly control error rates given individual differences across people and the various

sources of noise in people's performance. To account for this key source of variability, we binned participants into sextiles according to the true error rate they experienced during the training phase of the task. One sextile emerged with a mean error rate of 31.66%, which was closest to the theoretically predicted optimal error rate of 31.73%. Notably, this group exhibited the greatest LIs, and showed significantly greater learning than all other bins except one (Fig. 3C; see Table 1 for statistics).

The theoretical framework we employed not only produces an optimal error rate but can also be used to make learning rate predictions given any particular error rate. Thus, to garner additional evidence for or against our hypothesis we correlated the model's learning predictions at each error rate with subjects' empirically observed learning indices (LIs). The theoretically predicted learning rate was significantly correlated with the true LIs (Pearson correlation: R=0.19, p=0.01; Spearman correlation: ρ=0.16, p=0.03; Fig. 3D).

Table T1. Comparisons of learning indices (LIs) to the optimal error rate sextile

Mean Error Rate (SEM)	t	df	p	d
13.57% (0.61%)	3.26	60	0.002**	0.83
23.09% (0.57%)	2.48	57	0.02*	0.65
31.66% (0.42%)	-----	-----	-----	-----
44.17% (0.62%)	1.46	63	0.15	0.37
52.01% (0.38%)	2.18	53	0.03*	0.59
61.73% (0.55%)	2.49	56	0.02*	0.66

*Note: all comparisons are relative to the third sextile, with error rate 31.66%. * indicates p<0.05; ** indicates p<0.01.*

Taken together, theoretical predictions of an optimal motor learning error rate of ~30% were supported in a large data set consisting of people performing, and improving at, a visuomotor skill task. These data are consistent with the theoretical framework proposed by Wilson et al. [9], here generalized to the case of motor learning.

Discussion

Here we employed a novel Pong-based motor skill task, and provided behavioral evidence supporting the theoretical prediction that the optimal error rate for learning a motor skill may be ~30%. By controlling human subjects' error rates during our Pong task, we could influence how well they learned throughout the course of training. Our theoretical predictions, which were closely based on recent work on optimal learning for binary classification tasks [9], were supported both by showing the strongest learning in the group with the near-optimal error rate and a significant between-subjects correlation between the model predictions and the empirical data across a range of individuals' actual error rates (Figure 3). We propose that if an error rate of approximately ~30% is maintained throughout training of a simple one-dimensional motor skill, one may expect to observe the greatest improvements in said motor skill. This supports the work of "error-less" learning paradigms, where high success rates are maintained throughout training and the difficulty of the task increases as the learner improves [15,16]. Given the link between our theoretical predictions and the measured behavioral results, we also speculate that principles of gradient descent could be helpful for understanding the underlying psychological and neural algorithms for motor skill learning.

Why does the error rate during training influence learning?

In the mathematical framework proposed by Wilson et al. [9] and extended here, the relationship between training difficulty and learning is a necessary consequence of tuning the parameters

responsible for generating responses that reduce errors via gradient descent. Such an approach leaves a key question unanswered: “How does training difficulty result in changes in learning rate?” Multiple psychological explanations could be useful for addressing this question.

First, the “Challenge Point” framework [8] proposes that an agent’s performance at a motor skill is a source of information for future learning and action attempts. Highly expected outcomes, such as success when the task is easy, yield very little information as the agent can fully predict the consequences of their action. However, highly unexpected outcomes, such as success when the task is extremely difficult, are rare but can provide a high degree of information. However, when facing high difficulty, the potentially large amount of information available to the learner risks being uninterpretable: Was success due to sheer luck and/or external factors? What aspects of this motor plan made it more successful than previous attempts? The optimal “challenge point” occurs at a critical difficulty where a considerable amount of information can be learned from success while also maintaining a moderate degree of difficulty, such that outcome information is interpretable. In other words, this requires balancing the expectation of success (which decreases with increased difficulty and should be low to maximize potential information) with the functional difficulty (which should also be low to maximize the amount of interpretable information learned from feedback). The balance of minimizing both the expectation of success and the functional difficulty, to optimize the amount of interpretable information available for learning, results in a fixed optimal error rate during training [9].

Another explanation focuses on the role of reward and motivation in error-based learning. Motivation to perform well in a task can drive increased sensitivity to errors [17], which in turn can prompt faster learning [18]. Additionally, rewarding feedback tends to increase motivation [19] and can accelerate motor learning [20]. Thus, the recent history of successes and failures

may prime an “optimal” degree of motivation, where errors occur in concert with rewarded trials, keeping the agent’s error sensitivity and motivation for success high. Put another way, if the task is too easy, learning would be minimal; but, if the task is too difficult, successes are too infrequent to generate the benefits to motivation and error sensitivity that are crucial for maintaining a high learning rate. The balance of these two opposing forces could also result in an optimal error rate for learning.

The mathematical framework employed here does not allow us to differentiate between these explanations, or others not mentioned here. The formation of internal models for motor learning requires obtaining information about how observed sensory outcomes relate to internally generated action plans [10, 11, 21], in line with the Challenge Point hypothesis [8]. But reward feedback has also been shown to enhance motor learning [20, 22, 23], supporting a motivational account as well. Both explanations are plausible and may reflect how different brain regions (e.g., the cerebellum and basal ganglia, respectively) contribute to the fine tuning of motor learning-related parameters [24].

Extensions and limitations of the current framework

In our experiment, we chose a simple motor task where motor errors are generally assumed to be normally distributed. However, in some tasks, such as interval timing [25], the distribution of errors may not be normal and may follow other known distributions. As noted by Wilson and colleagues [9], the current approach still applies as long as performance improvements in those tasks are subserved by a gradient descent algorithm, but it would produce a different optimal error rate specific to that distribution of errors. In other words, an optimal error rate using non-Gaussian error distributions could be computed using the same equations described here. In tasks where the error distribution is unknown or cannot be mathematically described, empirical

observation or alternative methods of estimating the distribution must be applied in order to compute the optimal error rate.

We also restricted our consideration to a simple one-dimensional motor skill task (errors were only in the vertical dimension). However, many motor tasks involve multiple dimensions (e.g., other spatial dimensions, force, timing, etc.) which must be concurrently trained. Future work must be done to extend this framework to these more complex, naturalistic tasks. An important question which must be considered is whether learning in each dimension is fully independent, or if learning in one dimension affects learning in the others; the answer to this question would allow us to better model the error reduction processes involved in learning complex motor skills when multiple dimensions must be concurrently controlled [5]. Additionally, distributions of errors in each dimension may not follow similar mathematical descriptions, requiring more complicated multi-dimensional integrals to optimize the error rate. Nonetheless, the general mathematical framework demonstrated here can in theory be applied to such multi-dimensional tasks.

Lastly, our task only examines motor learning as defined as the gradual reduction in errors over time through practice, which arguably represents only one component of skill learning (i.e., improvements in “motor acuity” [26]). Explicit strategies and changes in one’s action policy, especially early in learning, can cause a learner to abruptly reduce their errors or discover new solutions [12, 27]. Explicit strategies and cognitive insights during skill learning likely arise from separate cognitive processes not considered here, though they still make important contributions to skill learning [28]-[30].

Future Applications

Our results demonstrate that the error rate during training affects the rate of improvement of a simple motor skill. We think this insight might be broadly applicable across applied domains. Rehabilitation paradigms, for example, have already recognized the need to adjust the difficulty of practice in order to optimally challenge patients and maximize the therapeutic benefit of training sessions [31, 32]. In sports coaching, the principle of progressively increasing the difficulty of practice has been widely applied and facilitates improvements in skill [33]-[35]. Optimistically, these approaches, in both sports pedagogy and rehabilitation, may benefit from the theory used here to mathematically uncover the sweet spot of training difficulty for maximum skill improvement.

Summary

We demonstrate that an optimal error rate for performance improvements in a motor skill may exist, and suggest that this optimal error rate is around ~30%. By testing a novel motor skill task (a modified version of Pong) on a large cohort of human participants, we provide empirical support for this prediction and demonstrate the potential importance of maintaining a particular error rate throughout training to maximize performance improvements. This approach may have broader implications in other domains concerned with optimizing the acquisition, refinement, and rehabilitation of motor skills.

Methods

Participants

200 participants (46% female; age: 27±5) were recruited online through Prolific. Recruitment was restricted to right-handed or ambidextrous individuals in the United States between the ages of 18 and 35, who had at least 40 prior Prolific submissions. All participants provided informed consent, in accordance with procedures approved by the Yale University Institutional

Review Board. Due to participant data not posting to the database, two participants did not have any data recorded. Four participants were further excluded due to not making any recorded movements during the task. In analyses of learning indices (see *Statistical Analyses*), subjects whose learning index deviated by 1.75 times the group interquartile range (IQR) from the median were excluded. This excluded 10 subjects. The final experimental sample was 184 subjects.

Task Design and Protocol

Participants were recruited on Prolific to play a modified version of the game Pong, written using PhaserJS 3 (phaser.io). Participants controlled a 15-pixel wide rectangular paddle, initially set to 200 pixels (px) in length and fixed 10 px from the left edge of the workspace, inside a 1000 px by 800 px arena enclosed within a white border. Movements of subjects' mouse or trackpad in the coronal plane directly translated to displacements of the paddle's vertical position (i.e., movements away from the body = up; toward = down). Participants were instructed to move the paddle to intercept a small circular puck (14 px diameter) which would move toward their paddle at an angle, bouncing off the walls of the arena. They were specifically instructed to try to cause the puck to hit the middle of the paddle, as this is the strategy most likely to result in success. This task required participants to integrate visual sensory information regarding the position of the puck into their movements of the paddle in order to successfully intercept the puck.

Each trial consisted of the puck crossing from 25 pixels left of the right side of the arena to the left side of the arena where the participant's paddle was located, crossing a horizontal distance of 975 px. At the start of each trial, the paddle's vertical position was always centered, and the puck began at a random vertical position sampled from a uniform distribution extending ±375

px from the center of the screen. The launch angle of the puck could be ±30-45°, to ensure that the trajectory of the puck was unpredictable. The puck began moving 1000ms after the start of the trial. If participants successfully intercepted the puck, a green “+10” text prompt would appear in the center of the screen for 1000 ms, indicating success. Otherwise, if the puck hit the edge of the arena without being intercepted by the paddle, a red “+0” text prompt was displayed in the center of the screen for 1000 ms, indicating failure. A running points total was displayed above the arena and participants were told to maximize the points they earned. After presentation of reward feedback, the paddle was re-centered, and the puck appeared in a new location to start the next trial.

Calibration and Error-control Algorithm

The experiment began with a 20-trial calibration phase, in which the horizontal speed of the puck was calibrated to be comfortable for each participant. The vertical speed of the puck was determined by the launch angle on a given trial. Since the duration of each trial was dependent on the horizontal speed of the puck, we needed to calibrate this speed to ensure every participant had ample time to react to the puck and adjust the paddle. The speed was calibrated using a staircase procedure: if the participant failed to intercept the puck, the horizontal speed was reduced by 75 px/s; if the participant successfully intercepted the puck, the speed increased by 75 px/s. The speed was initialized at 800px/s. The minimum allowed speed was 100px/s, but this was never achieved by any participant. This procedure resulted in a 72% (SD: 12%) success rate during the calibration phase.

Since this calibration phase often ended with a speed that slightly exceeded the comfortable range for the participant, the speed was cut to 80% of the final speed achieved. At this point, a 10-trial baseline began in which the puck speed and the paddle size were held constant. This

allowed us to assess each participants' initial degree of skill in the game. Mean success was 77% (SD:17%). At this point, the final 70-trial training phase began, and the paddle size was allowed to dynamically update according to a custom algorithm that monitored the recent distribution of errors to constrain participants' error rates. Each participant was randomly assigned to one of four groups, each with their own intended error rate: 15%, 30%, 50%, and 65%. On every trial, the puck's final vertical position relative to the final vertical position of the center of the paddle was added to a sliding 20-trial buffer, regardless of whether the trial was successful or not. After excluding outliers (defined as 3 times the median absolute deviation from the median; this is identical to the MATLAB *rmoutliers* function), we computed the root-mean-squared-error (RMSE) of this 20-trial history. From this we inferred a zero-mean Gaussian distribution with the RMSE as the standard deviation and, using the CDF of the Gaussian distribution, we computed a paddle size that would result in the intended error rate. This outlier criterion was chosen to minimize the number of excluded trials and only detect extreme outliers that could vastly impact the RMSE calculation for the duration of the 20-trial buffer, as determined by pilot experiments.

Statistical Analysis

The primary dependent measure on each trial was the task error, defined as the difference between the final vertical position of the puck and the final vertical position of the center of the paddle. Errors from every ten trials were binned in order to compute the RMSE, and these RMSEs were inverted in order to estimate precision, β. We note that the overall pattern of results does not change if the inverse of the standard deviation of the errors, an alternative measure of precision, was used instead of the inverse RMSE. RMSE was preferred since participants were instructed to aim to intercept the puck in the middle of the paddle, so any biases were random and unintentional. Each subject's baseline (trials 21-30) precision was subtracted from every

subsequent bin in the training phase in order to quantify improvement. Average improvement, or the learning index (LI), was quantified as the average β of all training bins after baseline correction. Statistical outliers, defined as 1.75 times the interquartile range (IQR) away from the median of each of the four error rate groups, were excluded. This excluded 10 people (5.15%), leaving us with a final experimental sample of 184.

Theoretical Modeling

Equation 4 gives an expression for the error rate (ER) as a function of the agent's precision β and the current error tolerance of the task Δ. To maximize the learning rate through gradient descent (Eq. 5), we must optimize $\frac{\partial ER}{\partial \beta}$, which, through differentiation is given by:

$$\frac{\partial ER}{\partial \beta} = 2\Delta \cdot p(\text{-}\beta\Delta) \tag{8}$$

Optimization of this expression requires setting the derivative with respect to Δ to zero:

$$0 = \frac{\partial}{\partial \Delta}\left(\frac{\partial ER}{\partial \beta}\right) = \frac{\partial}{\partial \Delta}\left(2\Delta \cdot p(\text{-}\beta\Delta)\right) \tag{9}$$

$$= 2p(\text{-}\beta\Delta) - 2\beta\Delta \cdot \frac{\partial}{\partial(\text{-}\beta\Delta)} p(\text{-}\beta\Delta)$$

This can be solved such that the optimal tolerance $\Delta^\star$ for a given β is given by

$$\beta\Delta^\star = \frac{p(\text{-}\beta\Delta^\star)}{p'(\text{-}\beta\Delta^\star)} \tag{10}$$

where $p(x)$ is the derivative of the distribution of errors with respect to the argument x. The optimal error tolerance $\Delta^\star$ changes with precision β, but the quantity $\beta\Delta^\star$ remains constant, implying that the optimal error rate $ER^\star$ from Equation 4 is constant, namely:

$$ER^\star = 2F(\text{-}\beta\Delta) \tag{11}$$

For a 1D motor task where errors are normally distributed, this optimal error rate is

$$ER^\star = 1 - erf\left(\frac{1}{\sqrt{2}}\right) \approx 0.3173 \tag{12}$$

The theoretical predictions plotted in Figure 1 and Figure 3D were computed in the manner described above. Since the integral of a Gaussian from $-\infty$ to x cannot be analytically solved, it was estimated using a 100,000-rectangle Reimann sum from -10 to x, a computationally costly but highly accurate estimate of the integral. 30,000 values of Δ from 0 to 3 in increments of 0.001 were used to generate the traces in figure 1 as well as to generate a function that translated Δ values into theoretically predicted learning rates $\frac{\partial ER}{\partial \beta}$, using the *approxfun* function in R [36]. Since the units of these predicted learning rates are arbitrary, we fit scaling and offset free parameters that allowed the predictions to match the observed precision improvements of our subjects using R's optim function most closely, and the resulting trace was plotted in Figure 3C.

References

1. Ohlsson S. Learning From Performance Errors. Psychol Rev. 1996;103(2):241–62.
2. Tseng Y, Diedrichsen J, Krakauer JW, Shadmehr R, Bastian AJ. Sensory Prediction Errors Drive Cerebellum-Dependent Adaptation of Reaching. J Neurophysiol. 2007 Jul;98(1):54–62.
3. Krakauer JW, Mazzoni P. Human sensorimotor learning: adaptation, skill, and beyond. Curr Opin Neurobiol. 2011 Aug 1;21(4):636–44.
4. Shmuelof L, Krakauer JW, Mazzoni P. How is a motor skill learned? Change and invariance at the levels of task success and trajectory control. J Neurophysiol. 2012 Jul 15;108(2):578–94.
5. Haith AM, Yang CS, Pakpoor J, Kita K. De novo motor learning of a bimanual control task over multiple days of practice. J Neurophysiol. 2022 Oct;128(4):982–93.
6. Yang CS, Cowan NJ, Haith AM. Control becomes habitual early on when learning a novel motor skill. J Neurophysiol. 2022 Nov;128(5):1278–91.
7. Sanger TD. Failure of Motor Learning for Large Initial Errors. Neural Comput. 2004 Sep;16(9):1873–86.
8. Guadagnoli MA, Lee TD. Challenge Point: A Framework for Conceptualizing the Effects of Various Practice Conditions in Motor Learning. J Mot Behav. 2004 Jun;36(2):212–24.
9. Wilson RC, Shenhav A, Straccia M, Cohen JD. The Eighty Five Percent Rule for optimal learning. Nat Commun. 2019 Nov 5;10(1):4646.
10. Pierella C, Casadio M, Mussa-Ivaldi FA, Solla SA. The dynamics of motor learning through the formation of internal models. PLOS Comput Biol. 2019 Dec 20;15(12):e1007118.
11. Thoroughman KA, Shadmehr R. Learning of Action Through Adaptive Combination of Motor Primitives. Nature. 2000 Oct 12;407(6805):742–7.
12. Taylor JA, Ivry RB. The role of strategies in motor learning. Ann N Y Acad Sci. 2012;1251(1):1–12.
13. Popa LS, Ebner TJ. Cerebellum, Predictions and Errors. Front Cell Neurosci [Internet]. 2019;12. Available from: https://www.frontiersin.org/articles/10.3389/fncel.2018.00524
14. Popa LS, Streng ML, Hewitt AL, Ebner TJ. The Errors of Our Ways: Understanding Error Representations in Cerebellar-Dependent Motor Learning. The Cerebellum. 2016 Apr 1;15(2):93–103.
15. Maxwell JP, Masters RSW, Kerr E, Weedon E. The implicit benefit of learning without errors. Q J Exp Psychol. 2001;54A(4):1049–68.
16. Capio CM, Poolton JM, Sit CHP, Holmstrom M, Masters RSW. Reducing errors benefits the field-based learning of a fundamental movement skill in children. Scand J Med Sci Sports. 2013;23(2):181–8.
17. Bengtsson SL, Lau HC, Passingham RE. Motivation to do Well Enhances Responses to Errors and Self-Monitoring. Cereb Cortex. 2009 Apr 1;19(4):797–804.
18. Wu SW, Trommershäuser J, Maloney LT, Landy MS. Limits to human movement planning in tasks with asymmetric gain landscapes. J Vis. 2006 Jan 23;6(1):5.
19. Bissonette GB, Roesch MR. Neurophysiology of Reward-Guided Behavior: Correlates Related to Predictions, Value, Motivation, Errors, Attention, and Action. In: Simpson EH, Balsam PD, editors. Behavioral Neuroscience of Motivation [Internet]. Cham: Springer International Publishing; 2016. p. 199–230. (Current Topics in Behavioral Neurosciences). Available from: https://doi.org/10.1007/7854_2015_382
20. Nikooyan AA, Ahmed AA. Reward feedback accelerates motor learning. J Neurophysiol. 2015 Jan 15;113(2):633–46.
21. Abdelghani MN, Lillicrap TP, Tweed DB. Sensitivity Derivatives for Flexible Sensorimotor

Learning. Neural Comput. 2008 Aug 1;20(8):2085–111.
22. Izawa J, Shadmehr R. Learning from Sensory and Reward Prediction Errors during Motor Adaptation. PLOS Comput Biol. 2011 Mar 10;7(3):e1002012.
23. Galea JM, Mallia E, Rothwell J, Diedrichsen J. The dissociable effects of punishment and reward on motor learning. Nat Neurosci. 2015 Apr;18(4):597–602.
24. Seidler RD, Kwak Y, Fling BW, Bernard JA. Neurocognitive Mechanisms of Error-Based Motor Learning. In: Richardson MJ, Riley MA, Shockley K, editors. Progress in Motor Control. New York, NY: Springer; 2013. p. 39–60. (Advances in Experimental Medicine and Biology).
25. Swanton DN, Matell MS. Stimulus Compounding in Interval Timing: The Modality–Duration Relationship of the Anchor Durations Results in Qualitatively Different Response Patterns to the Compound Cue. J Exp Psychol Anim Behav Process. 37(1):94–107.
26. Krakauer J, Hadjiosif A, Xu J, Wong A, Haith A. Motor Learning. In: Comprehensive Physiology. 2019. p. 613–63.
27. Mazzoni P, Krakauer JW. An implicit plan overrides an explicit strategy during visuomotor adaptation. J Neurosci Off J Soc Neurosci. 2006 Apr 5;26(14):3642–5.
28. McDougle SD, Ivry RB, Taylor JA. Taking Aim at the Cognitive Side of Learning in Sensorimotor Adaptation Tasks. Trends Cogn Sci. 2016 Jul 1;20(7):535–44.
29. McDougle SD, Wilterson SA, Turk-Browne NB, Taylor JA. Revisiting the Role of the Medial Temporal Lobe in Motor Learning. J Cogn Neurosci. 2022 Feb 1;34(3):532–49.
30. Haith AM, Krakauer JW. Model-Based and Model-Free Mechanisms of Human Motor Learning. In: Richardson MJ, Riley MA, Shockley K, editors. Progress in Motor Control. New York, NY: Springer; 2013. p. 1–21. (Advances in Experimental Medicine and Biology).
31. Metzger JC, Lambercy O, Califfi A, Dinacci D, Petrillo C, Rossi P, et al. Assessment-driven selection and adaptation of exercise difficulty in robot-assisted therapy: a pilot study with a hand rehabilitation robot. J NeuroEngineering Rehabil. 2014 Nov 15;11(1):154.
32. Lee H, Choi Y, Eizad A, Song WK, Kim KJ, Yoon J. A Machine Learning-Based Initial Difficulty Level Adjustment Method for Balance Exercise on a Trunk Rehabilitation Robot. IEEE Trans Neural Syst Rehabil Eng. 2023;31:1857–66.
33. Wickens CD, Hutchins S, Carolan T, Cumming J. Effectiveness of part-task training and increasing-difficulty training strategies: a meta-analysis approach. Hum Factors. 2013;55(2):461–70.
34. Zahran L, El-Beltagy M, Saleh M. A Conceptual Framework for the Generation of Adaptive Training Plans in Sports Coaching. In: Hassanien AE, Shaalan K, Tolba MF, editors. Proceedings of the International Conference on Advanced Intelligent Systems and Informatics 2019. Cham: Springer International Publishing; 2020. p. 673–84. (Advances in Intelligent Systems and Computing).
35. Marraffino MD, Schroeder BL, Fraulini NW, Van Buskirk WL, Johnson CI. Adapting training in real time: An empirical test of adaptive difficulty schedules. Mil Psychol. 2021 May 4;33(3):136–51.
36. Becker R. The new S language. CRC Press; 2018.