Title: Motor skill learning decreases movement variability and increases planning horizon

Authors

Luke Bashford[1,2,3,4,*], Dmitry Kobak[1,2,3,5,6,*], Jörn Diedrichsen[7], Carsten Mehring[1,2,3]

[1]Bernstein Center Freiburg, University of Freiburg, Freiburg, Germany

[2]Faculty of Biology, University of Freiburg, Freiburg, Germany

[3]Imperial College London, London, UK

[4]California Institute of Technology, Pasadena, CA, USA

[5]Champalimaud Centre for the Unknown, Lisbon, Portugal

[6]Institute for Ophthalmic Research, Tübingen University, Tübingen, Germany

[7]Brain and Mind Institute & Department for Computer Science, University of Western Ontario, Ontario, Canada

*****Equal contribution

Corresponding author

Luke Bashford, bashford@caltech.edu

Abstract

We investigated motor skill learning using a path tracking task, where human subjects had to track various curved paths at a constant speed while maintaining the cursor within the path width. Subjects' accuracy increased with practice, even when tracking novel untrained paths. Using a "searchlight" paradigm, where only a short segment of the path ahead of the cursor was shown, we found that subjects with a higher tracking skill differed from the novice subjects in two respects. First, they had lower movement variability, in agreement with previous findings. Second, they took a longer section of the future path into account when performing the task, i.e. had a longer planning horizon. We estimate that between one third and one half of the performance increase in the expert group was due to the longer planning horizon. An optimal control model with a fixed horizon (receding horizon control) that increases with tracking skill quantitatively captured the subjects' movement behaviour. These findings demonstrate that human subjects not only increase their motor acuity but also their planning horizon when acquiring a motor skill.

New and Noteworthy

We show that when learning a motor skill humans are using information about the environment from an increasingly longer amount of the movement path ahead to improve performance. Crucial features of the behavioural performance can be captured by modeling the behavioural data with a receding horizon optimal control model.

Introduction

The human motor system can acquire a remarkable array of motor skills. Informally, a person is said to be "skilled" if he or she can perform faster and at the same time more accurate movements than other, unskilled, individuals. What we don't know, however, is what learning processes and components underlie our ability to move better and faster. One component may be relatively "cognitive", involving the faster and more appropriate selection and planning of upcoming actions (Diedrichsen and Kornysheva 2015; Wong et al. 2015). Another component may be related to motor execution – the ability to produce and finely control difficult combinations of muscle activations, also called "motor acuity" (Shmuelof et al. 2012; Waters-Metenier et al. 2014). Depending on the structure of the task, changes in visuo-motor processing or feedback control may also contribute to skill development. Motor adaptation extensively studied using visuomotor and force perturbations (Shadmehr et al. 2010), may play a certain role in stabilizing performance, but it cannot by itself lead to improvements in the speed-accuracy trade-off (Wolpert et al. 2011).

A task commonly used in the experiments on motor skill learning is sequential finger tapping, where subjects are asked to repeat a certain tapping sequence as fast and as accurately as possible (Karni et al. 1995, 1998; Petersen et al. 1998; Walker et al. 2002). Improvement in such a task can continue over days, but previous papers have focussed mostly on the learning that is specific to the trained sequence(s) (Karni et al. 1995).

Many real-world tasks, however, do not involve the production of a fixed sequence of motor commands, but the flexible planning and execution of movements. Such flexibility is often well described by optimal feedback control models (Braun et al. 2009; Diedrichsen et al.

2010; Todorov and Jordan 2002) where the skilled actor appears to compute "on the fly" the most appropriate motor command for the task at hand. This requires demanding computations (Todorov and Jordan 2002), and the human motor system likely has found heuristics to deal with this complexity. One way to reduce complexity of the control problem is to not optimize the whole sequence of motor commands that will achieve the ultimate goal, but to only optimise the current motor command for a short distance into the future. This idea is called receding horizon control, also known as model predictive control (Kwon and Han 2005). Under this control regime, the system computes a feedback control policy that is optimal for a finite planning horizon. The control policy is then continuously updated as the movement goes on and the planning horizon is being shifted forward. This allows for adaptability, e.g. it can flexibly react to perturbations or unexpected challenges, as sensory information becomes available. Recent studies provided indirect evidence that favour the optimisation of short time-periods of a motor command (Dimitriou et al. 2013). The notion of planning horizon also arises in reinforcement learning, e.g. in the context of the so-called successor representation (Momennejad et al. 2017).

Motivated by these ideas, we propose that some of the skill of a down-hill skier or a race-car driver may lie not only in the increased ability to execute difficult motor commands (e.g. due to increased motor acuity), but also in the ability to plan further ahead and to optimize the movements for a longer time period into the future. In addition, we propose that the time span that subjects plan ahead increases with experience, leading to an increasing performance with training.

To test this idea, we designed an experimental condition which would allow us to measure the planning horizon that skilled actors are using when executing long sequence of movements that need to be planned “on the fly” – i.e. where the actual sequence of movements cannot be memorized. For this, we developed a path tracking task, where subjects had to maintain their cursor within a path that was moving towards them at a fixed speed. A similar task has been previously used in motor control research (Poulton 1974), using a mechanical apparatus with paths drawn on a paper roll that was moving at a fixed speed. It has been shown that subjects are able to increase their accuracy with training, but the different computational strategies between expert subjects and naïve performers remain unclear. In our study we use ‘searchlight’ trials in which subjects see various lengths of the approaching path ahead of their cursor to probe subjects forward planning and compare experts and novices in this respect.

Materials and Methods

Subjects

62 experimentally naïve subjects took part in this experiment (33 males and 29 females, age range 20-52 years old). Subjects gave written informed consent and were paid 10 €/h. The experimental procedures received ethics approval from the University of Freiburg.

Setup

Subjects sat at a desk looking at a computer monitor (Samsung Syncmaster 226BW) located ~80cm away. A cursor displayed on the screen (Matlab and Psychophysics Toolbox Version 3 (Brainard 1997)) was under position control by movements of a computer mouse. The mouse could be moved on the desk in all directions but only the horizontal (left and right) component contributed to the cursor movement: the vertical position of the cursor was fixed at 5.7mm above the base of the screen.

Task

To begin each trial subjects had to press the space bar. This displayed the cursor (R=2.9mm, 1.1cm from the bottom of the screen) and the path (width = 2.83cm) that extended from the top to bottom of the screen (30cm). The path continuously moved downward on the screen at a vertical speed of 34.1cm/s. The initially visible path was a straight line centered in the middle of the screen with the cursor positioned in the middle of the path. Once this initial section moved through the screen, the path then followed a random curvature (Fig. 1A). Subjects were instructed to keep the cursor between the path borders at all times moving only in the horizontal plane and were told to be as accurate as possible. The cursor and path were displayed in white if the cursor was within the path and both turned red when it was outside the path, always on a black background.

The cursor position was sampled at 60 Hz and the tracking accuracy was defined for each trial as the percentage of time steps when the cursor was inside the path. Running accuracy values were continuously displayed in the top left corner of the screen and final accuracies were displayed between the trials.

This experiment is based on a previous version where subjects were asked to track static randomly curved paths in 2D as quickly as possible without touching the sides [unpublished data, (Bashford et al. 2014)]. We later found that the 1D paradigm presented here was better suited to study the planning horizon as the speed was fixed.

Paradigm

Subjects were randomly assigned into two groups: expert (N=32) and naive (N=30). The paradigm included a training (expert group only) and a testing (all subjects) phase. Subjects in the expert group trained over 5 consecutive days, each day completing 30 minutes of path tracking (10 of 3-minute trials with short breaks in-between, searchlight length (s) 100%). If the performance improved from one trial to the next subjects saw a message saying "Congratulations! You got better! Keep it up!", otherwise the message "You were worse this time! Try to beat your score!" was shown. The training paths were randomly generated on the fly. Experts performed the testing set of trials after a short break following training on the final (5th) day. Naïve subjects performed only the testing set of trials.

The testing phase lasted 30 min (30 of 1-minute trials with breaks in-between) using 30 different pre-generated paths that were the same for all subjects. The testing phase in this experiment contained 3 normal trials (s=100%) and 27 searchlight trials (s=10-90%) where some upper part of the path was not visible. Three blocks of 10 trials with the searchlight length ranging from s=10% to s=100% (in steps of 10%) were presented, with the order shuffled in each block; the same fixed pseudorandom sequence was used for all subjects.

Path generation

Paths were generated before each trial start during training and a pre-generated fixed set was produced in the same way for testing. Each path was initialized to start at the bottom middle of the screen and the initial 30 cm of each path were following a straight vertical line. Subsequent points of the path midline had a fixed Y step of 40 pixels (1.1 cm) and random independent and identically distributed (iid) X steps drawn from a uniform distribution from 1 to 80 pixels (2.7mm – 2.2cm). Any step that would cause the path to go beyond the right or left screen edges was recalculated. The midline was then smoothed with a Savitzky-Golay filter (12th order, window size 41) and used to display path boundaries throughout the trial. All of the above parameters were determined in pilot experiments to create paths which were very hard but not impossible to complete after training.

Statistical analysis

In all cases, we used nonparametric rank-based statistical tests to avoid relying on the normality assumption. In particular, we used Spearman's correlation coefficient instead of the Pearson's coefficient, Wilcoxon signed-rank test instead of paired two-sample t-test, and Wilcoxon-Mann-Whitney rank sum test instead of unpaired two-sample t-test.

We initially recorded N=10 subjects in each group and observed statistically significant ($p<0.05$) effect that we are reporting here: positive correlation between the asymptote performance and the horizon length, as estimated via the changepoint and exponential models. We then recorded another N=20/22 (naïve/expert) subjects per group to confirm this finding. This internal replication confirmed the effect ($p<0.05$). The final analysis reported in this study was based on all N=62 subjects together. A preliminary version of the analysis for the initial N=10/10 subjects can be found in our preprint (Bashford et al. 2014), but note that it used a different way to estimate planning horizon compared to the procedure presented here, and so the values are not directly comparable.

Changepoint and Exponential model

We used two alternative models to describe the relationship between the searchlight length and the accuracy: a linear changepoint model and an exponential model. We used two different models to increase the robustness of our analysis and both models support our conclusions.

The changepoint model is defined by

$$y = \begin{cases} cs + o & \text{if } s \leq h_{cp} \\ ch_{cp} + o & \text{if } s > h_{cp} \end{cases}$$

where y is the subject's performance, s the searchlight length and (c, o, h_{cp}) are the subject-specific parameters of the model which define the baseline performance at searchlight 0% (o), the amount of increase of performance with increasing searchlight (c) and the planning horizon (h_{cp}) after which the performance does not increase any further.

The exponential model is defined by

$$y = \psi - \exp(-\rho s + d)$$

where the subject-specific parameters $(\psi,\ d, \rho)$ specify the performance at searchlight 0% $(\psi - \exp[d])$, the asymptote for large searchlights (ψ) and the speed of performance increase (ρ).

This function monotonically increases but it never plateaus. The speed of the increase depends on the parameter ρ with larger values meaning faster approaching the asymptote. We used the following quantity as a proxy for the “effective” planning horizon: $10+\log(5)/\rho$. It can be understood as the searchlight length that leads to performance being five times closer to the asymptote than at s=10%. The log(5) factor was chosen to yield horizon values of roughly the same scale as with the changepoint model above.

Both models (changepoint and exponential) were fit to the raw performance data of each subject, i.e. to the 30 data points, 3 for each of the 10 searchlight length values. The exponential fit (see Equation 2 in the Results) was done with the Matlab's nlinfit() function, implementing Levenberg-Marquardt nonlinear least squares algorithm. The changepoint fit (see Equation 1 in the Results) was done with a custom script that worked as follows. It tried all values of h_{cp} on a grid that included s=10% and then went from s=20% to s=100% in 100 regular steps. For each value of h_{cp} the other two parameters can be found via linear regression after replacing all $s>h_{cp}$ values with h_{cp}. We then chose h_{cp} that led to the smallest squared error.

Trajectory analysis

To shed light on the learning process we analysed additional parameters of the subjects’ movement trajectories.

First, we computed the time lag between the subjects' movement trajectories and the midline of the paths (Fig. 3A-B). To compute the lags, we interpolated both cursor trajectories and path midlines 10-fold (to increase the resolution of our lag estimates) and concatenated all three trials from the same subject and searchlight length. We computed the Pearson correlation coefficient between cursor trajectory and path midline for time shifts from of -300 to 300 ms, and defined the time lag as the time shift maximizing the correlation. We then used the obtained lags to compute mean-squared-error between the lagged path midline and the subject's trajectory for each subject and searchlight length (Fig. 3C-D).

Second, we extracted the cursor trajectories in all sections across all paths that shared a similar curved shape to explore the differences in cursor position at the apex of the curve (Fig. 4). The segments were selected automatically by sliding a window of length 18 cm across the path. We included all segments that were lying entirely to one side (left or right) of the point in the middle of the sliding window ("C-shaped" segments), with the upper part and the lower part both going at least 4.5 cm away in the lateral direction (see Fig. 4). Our results were not sensitive to modifying the exact inclusion criteria.

To draw the 75% coverage areas of the path inflection points in each group (Fig. 4), we first performed a kernel density estimate of these points using the Matlab function kde2d(), which implements an adaptive algorithm suggested in (Botev et al. 2010). After obtaining the 2d probability density function $p(x)$, we found the largest h such that $\int p(x)dx>0.75$ over the area where $p(x)>h$. We then used Matlab's contour() function to draw contour lines of height h in the $p(x)$ function.

Receding horizon model

We modeled subjects' behaviour by a stochastic receding horizon model in discrete time *t*. In receding horizon control (RHC,(Kwon and Han 2005)) motor commands u_t are computed to minimize a cost function L_t over a finite time horizon of length *h*:

$$\text{minimize } L_t(\{x_t\},\{u_t\})\#(1)$$

$$\text{subject to } L_t = \sum_{k=1}^{h} l_{t+k}$$

$$x_{t+1} = f(x_t, u_t)$$

where f defines the dynamics of the controlled system. Equation (1) is equivalent to an optimal control problem over the fixed future interval $[t+1, t+h]$. Solving (1) yields a sequence of optimal motor commands $\{u_0^{opt}, u_1^{opt}, \dots, u_{h-1}^{opt}\}$. The control applied at time *t* is the first element of this sequence, i.e. $u_t = u_0^{opt}$. Then, the new state of the system x_{t+1} is measured (or estimated) and the above optimization procedure is repeated, this time over the future interval $[t+2, t+1+h]$, starting from the state x_{t+1}.

Applying RHC to our experimental task, the dynamics of the cursor movement was modeled by a linear first-order difference equation:

$$x_{t+1} = x_t + u_{t-\tau} + \eta_t \quad \eta_t \in \mathcal{N}(0,\sigma^2)\#(2)$$

where *t* is the time step, x_t the cursor position at time *t*, u_t is the motor command applied at time *t* and τ the motor delay. η_t is the motor noise which was modeled as additive Gaussian white noise with zero mean and variance σ^2. We assumed that the controller minimizes the following cost function

$$L_t = \sum_{k=\tau+1}^{h} [-\log(q_{t+k}) + \lambda |u_{t-\tau+k-1}|^2]\#(3)$$

where L_t is the expected cost at time t, q_{t+k} is the probability of the cursor being inside the path at time $t+k$, h is the length of the horizon in time and λ is the weight of the motor command penalty. At every time step t, L_t is minimized to compute u_t. The cost function in (3) reflects a trade-off between accuracy (first term, i.e. $\log[q_{t+k}]$) and effort (second term) whereas their relative importance is controlled by λ. Cost functions with a similar accuracy-effort trade-off have been used previously to successfully model human motor behaviour (Braun et al. 2009; Diedrichsen 2007; Todorov and Jordan 2002).

We assumed that subjects have acquired a forward model of the control problem including the variance of the motor noise σ^2. We also assumed that subjects have an accurate estimate of the position of the cursor at time t, i.e. x_t is known. Under these assumptions the probability distribution of the cursor position at future times $t+k$, can be computed by:

$$p\left(x_{t+k} \middle| x_t, \{u_{t-\tau}, u_{t-\tau+1}, \dots, u_{t-\tau+k-1}\}\right) = \frac{1}{\sqrt{2\pi k\sigma^2}} e^{-\frac{(\hat{x}_{t+k})^2}{2k\sigma^2}} \#(5)$$

with

$$\hat{x}_{t+i} = x_t + \sum_{l=1}^{i} u_{t-\tau+l-1} \#(6)$$

The probability of the cursor being inside the path is then given by

$$q_{t+k} = \int_{m_{t+k}-\frac{w}{2}}^{m_{t+k}+\frac{w}{2}} \frac{1}{\sqrt{2\pi k\sigma^2}} e^{-\frac{(\hat{x}_{t+k}-z)^2}{2k\sigma^2}} dz \#(7)$$

where m_t is the position of the midline of the path at time t and w the width of the path. The receding horizon model assumes that motor commands u_t are computed by minimizing the cost L_t in each time step t for a fixed and known set of model parameters $(h, \lambda, \tau, \sigma^2)$. We simplify the optimisation problem by approximating q_{t+k} by

$$q_{t+k} \approx w \frac{1}{\sqrt{2\pi k\sigma^2}} e^{-\frac{(\hat{x}_{t+k}-m_{t+k})^2}{2k\sigma^2}} \#(8)$$

The higher $k\sigma_k^2$ is relative to the path width w, the higher the accuracy of this approximation. Note that the squared error is scaled by $k\sigma^2$ and hence, errors in the future are discounted. This is a consequence of the model of the cursor dynamics in equation (2).

Using equation (8) and removing all terms which do not depend on u_t, we can derive a simplified cost function

$$\tilde{L}_t = \sum_{k=\tau+1}^{h} \left[\frac{(\hat{x}_{t+k} - m_{t+k})^2}{2k\sigma^2} + \lambda |u_{t-\tau+k-1}|^2 \right] \#(9)$$

Equation (9) shows that the trade-off between accuracy and the magnitude of the motor commands is controlled by $\sigma^2\lambda$. We therefore can eliminate one parameter and use the equivalent cost function

$$\tilde{L}_t = \sum_{k=\tau+1}^{h} \left[\frac{(\hat{x}_{t+k} - m_{t+k})^2}{2k} + \tilde{\lambda} |u_{t-\tau+k-1}|^2 \right] \text{ with } \tilde{\lambda} = \sigma^2\lambda \#(10)$$

The gradient of $\tilde{L}_t$ is given by

$$\frac{\partial \tilde{L}_t}{\partial u_{t+j}} = 2\tilde{\lambda} u_{t+j} + \sum_{k=j+(\tau+1)}^{h} \left[\frac{(\hat{x}_{t+k} - m_{t+k})}{k} \right] \#(11)$$

with $j = 0, \ldots, h - (\tau + 1)$. The Hessian of $\tilde{L}_t$ is given by

$$\frac{\partial^2 \tilde{L}_t}{\partial u_{t+m} \partial u_{t+n}} = 2\delta_{m,n}\tilde{\lambda} + \sum_{k=\max(m,n)+(\tau+1)}^{h} \frac{1}{k} \#(12)$$

with m, n = $0, \ldots, h - (\tau + 1)$. For $\tilde{\lambda} = 0$ all pivots of the Hessian matrix in (12) are positive and therefore the Hessian is positive definite for $\tilde{\lambda} = 0$. For the general case $\tilde{\lambda} > 0$ the Hessian in (12) remains positive definite as $H_2 = H_1 + D$ is positive definite if H_1 is positive definite and D is a diagonal matrix with only positive diagonal entries. Given the positive definiteness of the Hessian in (12) we can conclude that the cost function $\tilde{L}_t$ is strictly convex with a unique global minimum. Setting the gradient (12) to **0** defines a system of $h-\tau$ linear equations with $h-\tau$ unknowns $(u_t, \ldots, u_{t+h-(\tau+1)})$ which solution minimizes $\tilde{L}_t$. The solution can be computed efficiently using standard numerical techniques. We used the 'linsolve' function of MATLAB which uses LU factorization.

When applying the model to the searchlight path we made the additional assumption that the model horizon increases with searchlight length s up to a maximal value h_{max} beyond which the model horizon remains constant:

$$h(s) = \begin{cases} s, & s < h_{max} \\ h_{max}, & s \geq h_{max} \end{cases} \#(14)$$

We used the same time step of 1/30s in the model as in the experiment. For a given set of model parameters $(h_{max}, \lambda, \tau, \sigma^2)$ we simulated the model 100 times with independent realizations of the motor noise. For each model trajectory we computed the time inside the path and the lag in the same way as they were computed for the subjects' trajectories. To obtain the time inside the path and the lag for a set of model parameters we averaged the obtained values across the 100 noise realizations.

The model was simulated on the searchlight paths to study the influence of the model horizon h_{max} and the motor noise σ^2 on performance and lag. To this end, we first simulated the model for the shortest searchlight paths (10%) assuming that h_{max} is at least as long as the searchlight length at 10% (=3cm) and using a motor delay of τ=200ms. The model was simulated using 50 logarithmically spaced values between 10^{-3} and 10^{+3} for λ and 45 values for σ^2 composed of 5 linearly spaced values between 0 and 0.04 and 40 linearly spaced values between 0.05 and 2. Together, this results in 45x50=2250 different parameter sets in total. From these sets, we chose values for λ and σ^2 which yielded a similar performance and lag as experimentally observed for the 10% searchlight (i.e. 45% time inside the path and a lag of 200ms). Using these parameter values, the model was then simulated for all searchlight paths for different model horizons. From the resulting performance as a function of searchlight lengths we computed the change-point in the same way as for the experimental data. In addition, the model was also simulated for different values of the motor noise and the change-point of the performance was computed for different noise levels as above. These analyses allowed us to investigate the influence of the model horizon and model motor noise on the change-point of the performance curve (see Fig. 5**)**. To establish the robustness of the model results, we repeated the above simulations and analyses for different values of the motor delay using τ=33ms, 100ms and 233ms.

Parts of the modeling computations were run on the high-performance computing cluster NEMO of the University of Freiburg (http://nemo.uni-freiburg.de) using Broadwell E5-2630v4 2.2 GHz CPUs.

All analysis code is available at https://github.com/dkobak/path-tracking.

Results

Learning the Tracking Skill

We designed an experiment where subjects had to a track a path moving towards them at a fixed speed (Fig. 1A and Methods). The narrow and wiggly path was moving downwards on a computer screen while the cursor had a fixed vertical position in the bottom of the screen and could only be moved left or right. Accuracy, our performance measure, was defined as the fraction of time that the cursor spent inside the path boundaries. One group of subjects (the expert group, N=32) trained this task for 30 minutes on each of 5 consecutive days. Another group (the naïve group, N=30) did not have any training at all. Both groups then performed a testing block that we describe below.

Over the course of five training days, the experts' accuracy increased from 66.9±8.0% to 79.6±6.4% (mean±SD across subjects, first and last training day respectively) as shown on Figs 1B-C, with the difference being easily noticeable and statistically significant ($p=8x10^{-7}$, z=4.9, Wilcoxon signed rank test; Cohen's d=1.8, N=32). As all paths generated during the training were different, this difference cannot be ascribed to memorizing the path, therefore this improvement represents the genuine acquisition of the skill of path tracking.

Searchlight testing

To unravel the mechanisms of skill acquisition we designed testing trials called "searchlight trials", during which subjects had to track curved paths as usual but could only see a certain part of the path (fixed distance s) ahead of the cursor. The searchlight length *s* varied between 10% and 100% of the whole path length in steps of 10% (the minimal s was ~3cm) to probe subjects' planning horizon. Searchlight testing was conducted after 5 days of training for experts or immediately for novices. During the testing block all subjects completed 30 one-minute-long trials (three repetitions of each of the 10 values of s). The average accuracy at full searchlight s=100% was 82.8±7.5% for the expert group and 65.7±8.4% for the naïve group (mean±SD across subjects), with the difference being highly significant ($p=2x10^{-9}$, z=6.0, Wilcoxon-Mann-Whitney rank sum test, Cohen's d=2.2, N=62). The performance of the naïve subjects during 100% searchlight trials (65.7±8.4%) was not significantly different from the initial performance of the expert subjects on their first day of training (66.9±8.0%), where searchlight was also 100% (p=0.76, z=0.3, Wilcoxon-Mann-Whitney rank sum test, Cohen's d=0.15, N=62).

Before we present the rest of the data, let us consider several possible ways in which the accuracy can depend on the searchlight length (Fig. 2A). For each subject, accuracy should be a non-decreasing function of searchlight length. The data presented in Poulton (1974) indicate that this function tends to become flat, i.e. subjects reach a performance plateau, after a certain value of the searchlight length that we will call *planning horizon* (Fig. 2A, top), while we assume all subjects will be constrained to the similar poor performance at the smallest searchlight. For the expert group, this function has to reach a higher point at s=100%, which could be achieved in one of two ways. Firstly, it could do so because the initial rise becomes steeper (Fig. 2A, bottom left), due to increased motor acuity after skill learning (Shmuelof et al. 2012, 2014). Alternatively, the expert group could reach a higher point at s=100% because the initial rise continues longer. This would suggest an increase in the planning horizon (Fig. 2A, bottom right) over which subjects plan and execute motor commands, described well by a receding horizon control (Kwon and Han 2005). It is likely a combination of both is employed by the human motor system during skill learning.

Fig. 2B shows subjects' accuracy in the searchlights trials as a function of the searchlight length s. All subjects were strongly handicapped at short searchlights, and at the shortest searchlight the performance of the two groups was similar with experts being only marginally better (42.5±2.3% for the expert group, 41.4±1.8% for the naïve group, $p=0.042$, $z=2.0$ Wilcoxon-Mann-Whitney rank sum test; Cohen's $d=0.5$, $N=62$).

Visual inspection of Fig. 2B suggests that both effects sketched in Fig. 2A contribute to expert performance. (i) the planning horizon for the expert group was longer than for the naïve group; and (ii) the expert group had higher accuracies in the initial part of the performance curve, before the performance plateaus, which could be explained by an increased motor acuity.

To investigate differences in tracking skill between groups, we estimated the planning horizons of individual subjects. For this we fit each subject's performance (y) with a changepoint linear-constant curve (see Methods), where the location of the changepoint defines the horizon length. The initial slope of the changepoint model was significantly different between the two groups (3.7±1.2 %/cm in the expert group vs. 3.0±1.2 %/cm in the naïve group, mean±SD; medians: 3.6 %/cm vs. 2.6 %/cm, $p=0.008$, $z=2.6$, Wilcoxon-Mann-Whitney rank sum test; Cohen's $d=0.6$, $N=62$). Fig. 2C shows that there was a positive correlation between the initial slope and asymptote accuracy ($R=0.49$, $p=6\times10^{-5}$ Spearman correlation, $N=62$).

At the same time, we found that the novice group had an average horizon length of 11.5±3.6cm (mean±SD; median: 12.0cm) and the expert group a horizon length of

14.2±3.5cm (median: 13.2cm), with statistically significant difference ($p=0.007$, $z=2.7$, Wilcoxon-Mann-Whitney rank sum test; Cohen's $d=0.8$, $N=62$). We also found a positive correlation between the horizon length and the asymptotic performance ($R=0.34$, $p=0.006$, Spearman correlation, $N=62$) (Fig. 2D).

In addition to the changepoint model, we also quantified the "effective" planning horizon using a single exponential to fit the individual subjects' performance data (see Methods). This analysis confirmed our results (Fig. 2E). We again observed a significant difference in the effective horizon length between the two groups (14.76±4.6cm vs. 11.04±4.7cm, means±SD for both groups, medians: 13.6cm and 10.7cm, $p=0.002$, $z=3.0$, Wilcoxon-Mann-Whitney rank sum test; Cohen's $d=0.8$, $N=62$). Again, we found a positive correlation between the asymptote performance and the effective horizon length ($R=0.43$, $p=0.0008$, Spearman correlation, $N=62$).

We therefore conclude that the difference between expert and naïve performances is a combination of both possibilities presented in Fig. 2A. Using the expert and naive median estimates of the intercept, the slope, and the horizon in the changepoint model, we can estimate the contribution of both effects on the asymptote performance. The changepoint model asymptote performance for the naive group was 63.5%, compared to 78.7% for the expert group. The model performance of the expert group at the naive horizon was 74.2%. Hence, approximately 71% of the expert performance gain of 15.2%, was due to the increase in the initial slope (possibly due to increased motor acuity), and the remaining 29% can be attributed to the increase in planning horizon. The identical procedure with mean model parameter estimates instead of median estimates, yields 44% attributable to motor acuity and

56% attributable to planning horizon. However, these results do not elucidate whether these processes are causally related (see Discussion).

Trajectory analysis

Naïve subjects performed worse than the expert subjects at long searchlights but all subjects performed almost equally badly at short searchlights. What kinematic features can these differences be attributed to?

Clearly, at short searchlights, performance has to be reactive. To measure how quickly changes in the path were reflected in the motor commands, we computed the time lag between cursor trajectory and path midline (the lag maximizing cross-correlation between them). This analysis was done by pooling all trials for each subject and searchlight length together (see Methods). As Fig. 3A shows, the lag was ~200 ms at s=10% for all subjects and dropped to ~0 ms at s=50% for the expert group. While many naïve subjects also decreased their lags to zero, 10 out of 30 never achieved the 0 ms lag. The five naïve subjects showing the largest lags at large searchlights were also those with the worst performance (Fig. 3B). Therefore, there was a strong negative correlation between the asymptote lag (mean across s=80-100%) and the asymptote performance (mean across s=80-100%) of R=-0.58 (Fig. 3B, $p=8x10^{-7}$ Spearman correlation, N=62).

We used the obtained lags to compute the root-mean-squared-error (RMSE) between the cursor trajectory and the lagged midline. The obtained RMSE was consistently lower in the expert group than in the naive group, with difference increasing with searchlight length (Fig. 3C). The asymptote RMSE was 1.67±0.31cm (mean±SD across subjects; median: 1.69) in the naive group and 1.23±0.30 cm (median: 1.13) in the expert group ($p=2.14x10^{-6}$, z=4.74, Wilcoxon-Mann-Whitney rank sum test; Cohen's d=1.44, N=62), and was negatively correlated with the asymptote performance (Fig. 3D; R=-0.79, p=0, Spearman correlation, N=62). This shows that the naive subjects were not simply lagging behind the optimal trajectory, but did larger errors even after accounting for the lag.

Next, for each testing path we found all segments exhibiting sharp leftward or rightward bends (see materials and methods, our inclusion criteria yielded 13±5 segments per path, mean±SD). For each searchlight length and for each subject, we computed the average cursor trajectory over all segments (N=38±8 segments per searchlight) after aligning all segments on the bend position (Fig. 4, leftward bends were flipped to align them with the rightward bends). At s=10% all subjects from both groups follow very similar lagged trajectories, resulting in low accuracy. As searchlight increases, expert subjects reach zero lag and choose more and more similar trajectories, whereas naïve subjects demonstrate a wide variety of trajectories with some of them failing to reach zero lag and others failing to keep the average trajectory inside the path boundaries. To visualize this, we plotted the kernel density estimate 75% coverage contour of inflection points for each group. As the searchlight increases, the groups become less overlapping and the naïve group appears to form a bimodal distribution (Fig. 4).

To study the changes in movement variability produced by subjects after skill learning, we additionally looked at the within-subject variability during the segments defined above at 100% searchlight and compared this across groups. To measure the variability in subject's movement on a single subject level, we summed the standard deviations in both x and y directions across inflection points of each single segment. The subjects' averages of these positions are shown in Fig. 4. The summed standard deviations were significantly lower for the expert group (median summed SD=24.1) than for the naïve group (median summed SD=37.6) (p=1.5x10-7, z=-5.2, Wilcoxon-Mann-Whitney rank sum test, Cohen's d=1.7, N=62). This effect was not present at 10% searchlight (expert median summed SD=53.3, naïve median summed=49.9, p=0.94, z=0.08, Wilcoxon-Mann-Whitney rank sum test, Cohen's d=0.03, N=62).

In summary, at very short searchlights all subjects performed poorly because in this reactive regime their trajectories lagged behind the path. At longer searchlights the expert subjects were able to plan their movement to accommodate the bends (the longer the searchlight the better), but naïve subjects failed to do so in various respects: either still lagging behind, not being able to execute the fine movements due to lower motor acuity and higher movement variability, or not being able to plan a good trajectory.

Receding horizon model analysis

Next, we modeled subjects' behaviour by receding horizon control (RHC). Previously, optimal feedback control models have been proposed as mechanisms by which the human brain computes motor commands. Here we intend to expand this framework to include receding horizon control, a version of optimal feedback control with finite horizons, as a mechanism by which motor commands are computed by the human brain. We illustrate this here by showing that such an approach is able to capture some crucial features of the behavioural results from our experiments.

In RHC, a sequence of motor commands is computed to minimize the expected cost over a future time interval of finite length, i.e. the horizon. After the first motor command is applied, the optimization procedure is repeated using a time interval shifted one time step ahead. See

Methods section for a more detailed and formal description of RHC. As cost function, we used the weighted sum of a measure of inaccuracy (i.e. probability of being outside the path) and the magnitude of the motor cost (see Methods for details). Cost functions with a similar trade-off between movement accuracy and motor command magnitude have been used previously to describe human motor behaviour in different tasks (Braun et al. 2009; Diedrichsen 2007; Todorov and Jordan 2002). The model has four different parameters: horizon (h_{max}), motor noise (σ^2), motor delay (τ) and motor command penalty weight (λ).

We ran the model on the experimental paths to obtain simulated movement trajectories from which task performance and lag could be computed in the same way as for the experimental trajectories (Fig. 2 and 3). To determine the model parameters λ and σ^2, we simulated the model for the shortest searchlight paths (10%) using different values of λ and σ^2 while fixing τ at 200ms and assuming that the horizon covers at least the length of the 10% searchlight (Fig. 5A-C). From this parameter scan, we determined the values of λ (0.776) and σ^2 (0.271) for which the model yielded approximately the experimentally observed task performance and lag of 45% and 200ms for the 10% searchlight paths (cf. Fig. 2 and 3).

Using these parameter values we then simulated the model for all searchlight paths for varying values of the horizon (Fig. 5D, E). As a sanity check, we also simulated the model for varying values of the motor noise with a fixed value of the horizon h_{max}=14.8cm (Fig. 5G, H). Our simulations revealed that both, a larger model horizon as well as a smaller motor noise parameter increased the task performance and decreased the lag for large searchlight lengths. Hence, the experimentally observed higher performances and smaller lags of expert subjects compared to naive (Fig. 2B and 3A) could be explained either by an increased model horizon or by reduced motor noise in the model. However, the searchlight length at which the task performance of the model reached a plateau increased with model horizon while it remained constant or decreased with a smaller motor noise parameter (Fig. 5F, I). Experimentally, on the other hand, we observed that subjects with a higher task performance reached their performance plateau at higher searchlights (Fig. 2D, E). This correlation between performance and plateau onset, that was observed experimentally, cannot be explained by the variation of the motor noise parameter across subjects, but is only consistent with an increase of the model horizon parameter for subjects with higher performance. Moreover, with changing motor noise, the model predicted substantial changes in task performance and lag not only for large but also for short searchlight lengths while experimentally the differences between expert and naïve subjects were small for the 10% searchlight length. Model predictions for changes in planning horizon were again consistent with this experimental observation.

The analyses of Fig. 5G-I were repeated for various model horizons between 3.4cm and 29.6cm (not shown) showing similar patterns as presented in Fig. 5G-I for h_{max} >3.4cm: the performance change-point tended to remain constant or increase with increasing motor noise; changing motor noise induced a clear change of performance and lag also at short searchlights. For h_{max}=3.4cm the performance was essentially constant across searchlights for all values of the noise. Furthermore, we repeated the model simulations for motor delays of τ=33ms, τ=100ms and τ=233ms and obtained qualitatively similar results (not shown).

Discussion

We used a paradigm that allowed us to study skill development when humans had to track an unpredictable spatial path. The skill requires fast reactions to new upcoming bends in the road, but also a substantial "planning ahead" component – i.e. the anticipation and preplanning of movements that have to be made in the near future. We used the accuracy, i.e. the fraction of time the cursor was inside the path boundaries, as the measure of performance. We observed a substantial improvement in accuracy after 5 days of training (Fig. 1B,C). The paths were different on every trial, so the improvement in performance cannot be attributed to a memory for the sequence.

What changes in the motor system occur through learning that allowed skilled subjects to perform better? One component of this improvement is motor acuity (Shmuelof et al. 2012, 2014) and corresponds to the subjects' ability to execute motor commands more accurately. We hypothesized that an additional component is an increased ability to take into account approaching path bends and to prepare for an upcoming movement segment. We directly estimated both effects by using a searchlight testing where only a part of the approaching curve was visible. In agreement with our hypothesis, we found that subjects with a higher tracking skill demonstrated larger planning horizons: on average ~14cm for the expert group vs. ~11cm for the naïve group, corresponding to the time horizons of ~0.4s and ~0.3s respectively. Our results suggest that the increase in planning horizon is not an epiphenomenon but is causally related to the performance increase, as expert subjects showed worse performance when the searchlight was reduced below their planning horizon (Fig. 2C). We estimate that between one third and one half of the performance increase in the expert group was due to the longer planning horizon.

The improvement of acuity and the extension of planning horizon are not necessarily independent processes and may influence each other. For example, it is possible that improved acuity frees up cognitive resources that allow the expansion of planning horizon. Future work should investigate the causal relationship between these two aspects of skill learning.

Note that "planning"/"preparing" the movement can be interpreted differently depending on the computational approach. In the framework of optimal control (Todorov and Jordan 2002), subjects do not plan the actual trajectory to be followed, but instead use an optimal time-dependent feedback policy and then execute the movement according to this policy. The observed increase in planning horizon can be interpreted in the framework of model predictive control, also known as receding horizon control, RHC (Kwon and Han 2005). In RHC, the optimal control policy is computed for a finite and limited planning horizon, which may not capture the whole duration of the trial. This policy is then applied for the next control step, which is typically very short, and the planning horizon is then shifted one step forward to compute a new policy. Hence, RHC does not use a pre-computed policy, optimal for an infinite horizon, but a policy which is only optimal for the current planning horizon. Increasing the length of the planning horizon is therefore likely to increase the accuracy of the control policy. In our experiments this would allow for a larger fraction of time spent within the path boundaries. We designed a simple RHC model to test directly which components in the model would have to change through training to quantitatively explain the subject's behaviour. The dynamics of movement and the cost function were modeled in line with previous studies that used optimal control to describe human behaviour in various motor control and learning tasks (Braun et al. 2009; Diedrichsen 2007; Todorov and Jordan 2002).

We ran our RHC model on the experimental paths and demonstrated that it yields qualitatively correct predictions: larger value of model horizons led to performance similar to that of the experts' subjects. Our findings, thus, demonstrate that subjects' behaviour can be understood in the context of RHC, and longer planning horizons of the expert group indicate that subjects learn how to take advantage of future path information to improve motor performance.

Despite a clear difference in the distribution of planning horizons between the naive and the expert groups (Fig. 2D), there was a substantial overlap: the planning horizon of many naive and expert subjects were similar. While this might simply reflect a moderate effect size combined with inter-subject variability and measurement noise, it also remains a possibility that the difference between groups was largely caused by those naive subjects with very low horizons and expert subjects with very high horizons.

Related work

Ideas like the RHC were put forward in a recent study (Ramkumar et al. 2016) that suggested that movements are broken up in 'chunks' in order to deal with the computational complexity of planning over long horizons. That study suggests that monkeys increase the length of their movement chunks during extended motor learning over the course of many days which may be explained by monkeys increasing their planning horizon with learning. At the same time, the efficiency of movement control within the chunks improved with learning which may also be the result of a longer horizon. Despite these potential consistencies with our approach we note that in their model Ramkumar et al. (2016) assumed that 'chunks' are separated by halting points (i.e. points of zero speed) and movements within 'chunks' are optimized independently from each other. Our RHC model does not have independent movement elements but movements are optimized continuously.

Even though our study, to the best of our knowledge, is the first to directly investigate the evolution of the planning horizon during continuous path tracking, an increase in the planning horizon after learning has been recently demonstrated when learning sequences of finger movements (Ariani et al. 2020). Similar path tracking tasks have been used before (Poulton 1974). Using a track that was drawn on a rotating paper roll, these early studies found that the accuracy of the tracking increased with practice and with increasing searchlight length (which was modified by physically occluding part of the paper roll, (Poulton 1974), p 187). These studies, however, did not investigate the effect of learning on the planning horizon.

More recent studies used path tracking tasks where the goal was to move as fast as possible while maintaining the accuracy (instead of moving at a fixed speed). In all of these studies the identical path was repeatedly presented. In one study subjects had to track a fixed maze without visual feedback and learnt to do it faster as the experiment progressed (Petersen et al. 1998); there the subjects had to once “discover” and then remember the correct way through the maze. In another series of experiments, Shmuelof et al. asked subjects to track two fixed semi-circular paths. Subjects became faster and more accurate over the course of several days (Shmuelof et al. 2012), but this increase in the speed and accuracy did not generalize to untrained paths (Shmuelof et al. 2014). In contrast to these previous path tracking studies, we used randomly generated paths throughout the experiment. By investigating the generalization of the path tracking skill to novel paths we could reveal an increasing planning horizon with learning.

The planning horizon could possibly be inferred from the distance the subject gazes ahead of the cursor, and the inclusion of eye movement data could provide an interesting extension for future work. However, eye movement recordings often do not provide a clear answer in these types of experiments (Green and Bavelier 2006; Lehtonen et al. 2014; Wilkie et al. 2008) and the searchlight paradigm remains the most direct avenue of establishing how much information is used to guide behaviour.

Conclusion

In conclusion, we have established that people are able to learn the skill of path tracking and improve their skill over 5 days of training. This increase in motor skill is associated with the increased motor acuity and increased planning horizon. The dynamics of preplanning can be well described by a receding horizon control model.

Acknowledgements

The study was in parts supported by the German Federal Ministry of Education and Research (BMBF) grant 01GQ0830 to BFNT Freiburg-Tübingen. The authors acknowledge support by the state of Baden-Württemberg through bwHPC and the German Research Foundation (DFG) through grant no INST 39/963-1 FUGG. The authors also thank the 'Struktur- und Innovationsfonds Baden-Württemberg (SI-BW)" of the state of Baden-Württemberg for funding.

Author Contribution

Conceptualization, LB. DK. and CM; Methodology, LB. DK. JD and CM; Formal Analysis, LB. DK and CM; Writing – Original Draft, LB. DK and CM; Writing – Review and Editing, LB, DK, JD and CM.

Figure Legends

Figure 1. Experimental Paradigm. (A) Subjects had to track a curved path that was dropping down from top to bottom of the screen with a fixed speed of 34 cm/sec by moving the cursor horizontally. (B) Expert subjects' performance over the 5 days of training. Bold line shows the group average, thin lines show individual subjects (each point is a mean over 3 trials with the same searchlight length, 100%). (C) Expert subjects' performance over the 5 days of training with the performance on the first day subtracted.

*Figure 2. Searchlight testing. (A) Expert subjects were trained to have a higher performance at full searchlight length (top). This could be achieved by an increased initial slope (bottom left) at smaller searchlight length and/or an increased planning horizon as indicated with dashed vertical lines (bottom right). (B) Mean tracking performance for each searchlight length for each individual subject, in blue for the expert group and in red for the naïve group. Faint lines show individual subjects and bold lines show group means. (C) Relationship between the asymptote performance and the initial slope in in the changepoint linear-constant model (***: $p<0.001$). (D-E) Planning horizon for each subject was defined by fitting a changepoint linear-constant curve (D) or an exponential curve (E) (see text). Both models yield an asymptote performance for each subject; the changepoint model yields a horizon length and the exponential fit yields an "effective" horizon length. The scatter plots with marginal distributions show relation between the asymptote performance (as a proxy for subjects' skill) and their planning horizon. Spearman's correlation coefficients are shown on the plot (**: $p<0.01$, ***: $p<0.001$). Colour of the dot indicates the group.*

*Figure 3. Analysis of trajectories. (A) Mean time lag between cursor trajectory and path midline, for each searchlight length for each individual subject (faint lines) and mean of per-subject values (bold lines), in blue for the expert group and in red for the naïve group. (B) Asymptote lag and asymptote performance across subjects. Correlation coefficient is shown on the plot (***p<0.001). Colour of the dot indicates the group. (C) and (D) show the same for the root mean square error (RMSE) between the cursor trajectory and the path midline.*

Figure 4. Average per-subject trajectories in sharp bends (leftward bends were flipped to align them with the rightward bends). Each trajectory is averaged across approximately 40 bends identified in all paths (the number of bends varied across searchlight lengths, see Methods section 'Trajectory analysis'). Colour of the lines indicates the group. Black lines show average path contour. Dots show turning points of the trajectory. Contour lines show the kernel density estimate 75% coverage areas. Subplots correspond to searchlight lengths s=10%, 20%, 50%, 60%, 90% and 100%.

Figure 5: Task performance (A) and lag (B) for model simulations of the 10% searchlight paths as a function of λ *and* σ^2 *assuming* τ*=200ms and a model horizon of at least the length of the 10% searchlight. Note that for high values of* λ *and* σ^2 *the lag may not be reliable due to small peaks in the cross-correlogram between the movement trajectories predicted by the model and the paths (C). The white lines in (A) and (B) show the value of the parameters which yielded a task performance and a lag similar to what was experimentally observed (Fig.2 and 3). The intersection of the two lines was at* $\lambda = 0.77$ *and* $\sigma^2 = 0.27$*. Using these parameter values the model was simulated for all searchlight lengths and for various model horizons yielding the task performance and lag shown in (D) and (E). The performance curves in (D) were analysed using the same change-point analysis as for the experimentally*

obtained performance curves demonstrating an increasing change-point with increasing model horizon before flattening out at around 10cm (F) except for very short model horizons for which the performance curves were essentially flat and therefore, the change-point could not reliably be detected. Using a fixed horizon of h_{max}*=14.8cm and* λ *=0.77 model performance and lag was computed for varying motor noise (G, H) and the change-point of the performance was calculated for different values of the motor noise (I).*

References

Ariani G, **Kordjazi N**, **Diedrichsen J**. The planning horizon for movement sequences. *bioRxiv* 2020.07.15.204529, 2020.

Bashford L, **Kobak D**, **Mehring C**. Motor skill learning by increasing the movement planning horizon [Online]. *ArXiv14106049 Q-Bio* , 2014http://arxiv.org/abs/1410.6049 [15 Dec. 2015].

Botev ZI, **Grotowski JF**, **Kroese DP**. Kernel density estimation via diffusion. *Ann Stat* 38: 2916–2957, 2010.

Brainard DH. The Psychophysics Toolbox. *Spat Vis* 10: 433–436, 1997.

Braun DA, **Aertsen A**, **Wolpert DM**, **Mehring C**. Learning Optimal Adaptation Strategies in Unpredictable Motor Tasks. *J Neurosci* 29: 6472–6478, 2009.

Diedrichsen J. Optimal task-dependent changes of bimanual feedback control and adaptation. *Curr Biol CB* 17: 1675–1679, 2007.

Diedrichsen J, **Kornysheva K**. Motor skill learning between selection and execution. *Trends Cogn Sci* 19: 227–233, 2015.

Diedrichsen J, **Shadmehr R**, **Ivry RB**. The coordination of movement: optimal feedback control and beyond. *Trends Cogn Sci* 14: 31–39, 2010.

Dimitriou M, **Wolpert DM**, **Franklin DW**. The Temporal Evolution of Feedback Gains Rapidly Update to Task Demands. *J Neurosci* 33: 10898–10909, 2013.

Gelman A, **Carlin JB**, **Stern HS**, **Rubin DB**. *Bayesian Data Analysis, Second Edition*. 2nd edition. Boca Raton, Fla: Chapman and Hall/CRC, 2003.

Green CS, **Bavelier D**. Effect of action video games on the spatial distribution of visuospatial attention. *J Exp Psychol Hum Percept Perform* 32: 1465–1478, 2006.

Karni A, **Meyer G**, **Jezzard P**, **Adams MM**, **Turner R**, **Ungerleider LG**. Functional MRI evidence for adult motor cortex plasticity during motor skill learning. *Nature* 377: 155–8, 1995.

Karni A, **Meyer G**, **Rey-Hipolito C**, **Jezzard P**, **Adams M**, **Turner R**, **Ungerleider L**. The acquisition of skilled motor performance : Fast and slow experience-driven changes in primary motor cortex. *Proc Natl Acad Sci U S A* 95: 861–868, 1998.

Kwon WH, **Han SH**. Receding Horizon Control: Model Predictive Control for State Models [Online]. Springer-Verlag.//www.springer.com/la/book/9781846280245 [21 Mar. 2018].

Lehtonen E, **Lappi O**, **Koirikivi I**, **Summala H**. Effect of driving experience on anticipatory look-ahead fixations in real curve driving. *Accid Anal Prev* 70: 195–208, 2014.

Momennejad I, **Russek EM**, **Cheong JH**, **Botvinick MM**, **Daw ND**, **Gershman SJ**. The successor representation in human reinforcement learning. *Nat Hum Behav* 1: 680–692, 2017.

Petersen SE, **Mier H van**, **Fiez JA**, **Raichle ME**. The effects of practice on the functional anatomy of task performance. *Proc Natl Acad Sci* 95: 853–860, 1998.

Poulton EC. *Tracking Skill and Manual Control*. Academic Press, Incorporated, New York, N.Y., 1974.

Ramkumar P, **Acuna DE**, **Berniker M**, **Grafton ST**, **Turner RS**, **Kording KP**. Chunking as the result of an efficiency computation trade-off. *Nat Commun* 7: 12176, 2016.

Shadmehr R, **Smith MA**, **Krakauer JW**. Error Correction, Sensory Prediction, and Adaptation in Motor Control. *Annu Rev Neurosci* 33: 89–108, 2010.

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* 578–594, 2012.
Shmuelof L, **Yang J**, **Caffo B**, **Mazzoni P**, **Krakauer JW**. The neural correlates of learned motor acuity. *J Neurophysiol* 112: 971–980, 2014.
Todorov E, **Jordan MI**. Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5: 1226–1235, 2002.
Walker MP, **Brakefield T**, **Morgan A**, **Hobson JA**, **Stickgold R**. Practice with Sleep Makes Perfect : Sleep-Dependent Motor Skill Learning. *Neuron* 35: 205–211, 2002.
Waters-Metenier S, **Husain M**, **Wiestler T**, **Diedrichsen J**. Bihemispheric Transcranial Direct Current Stimulation Enhances Effector-Independent Representations of Motor Synergy and Sequence Learning. *J Neurosci* 34: 1037–1050, 2014.
Wilkie RM, **Wann JP**, **Allison RS**. Active gaze, visual look-ahead, and locomotor control. *J Exp Psychol Hum Percept Perform* 34: 1150–1164, 2008.
Wolpert DM, **Diedrichsen J**, **Flanagan JR**. Principles of sensorimotor learning. *Nat Rev Neurosci* 12: 739–51, 2011.
Wong AL, **Lindquist MA**, **Haith AM**, **Krakauer JW**. Explicit knowledge enhances motor vigor and performance: motivation versus practice in sequence tasks. *J Neurophysiol* 114: 219–232, 2015.

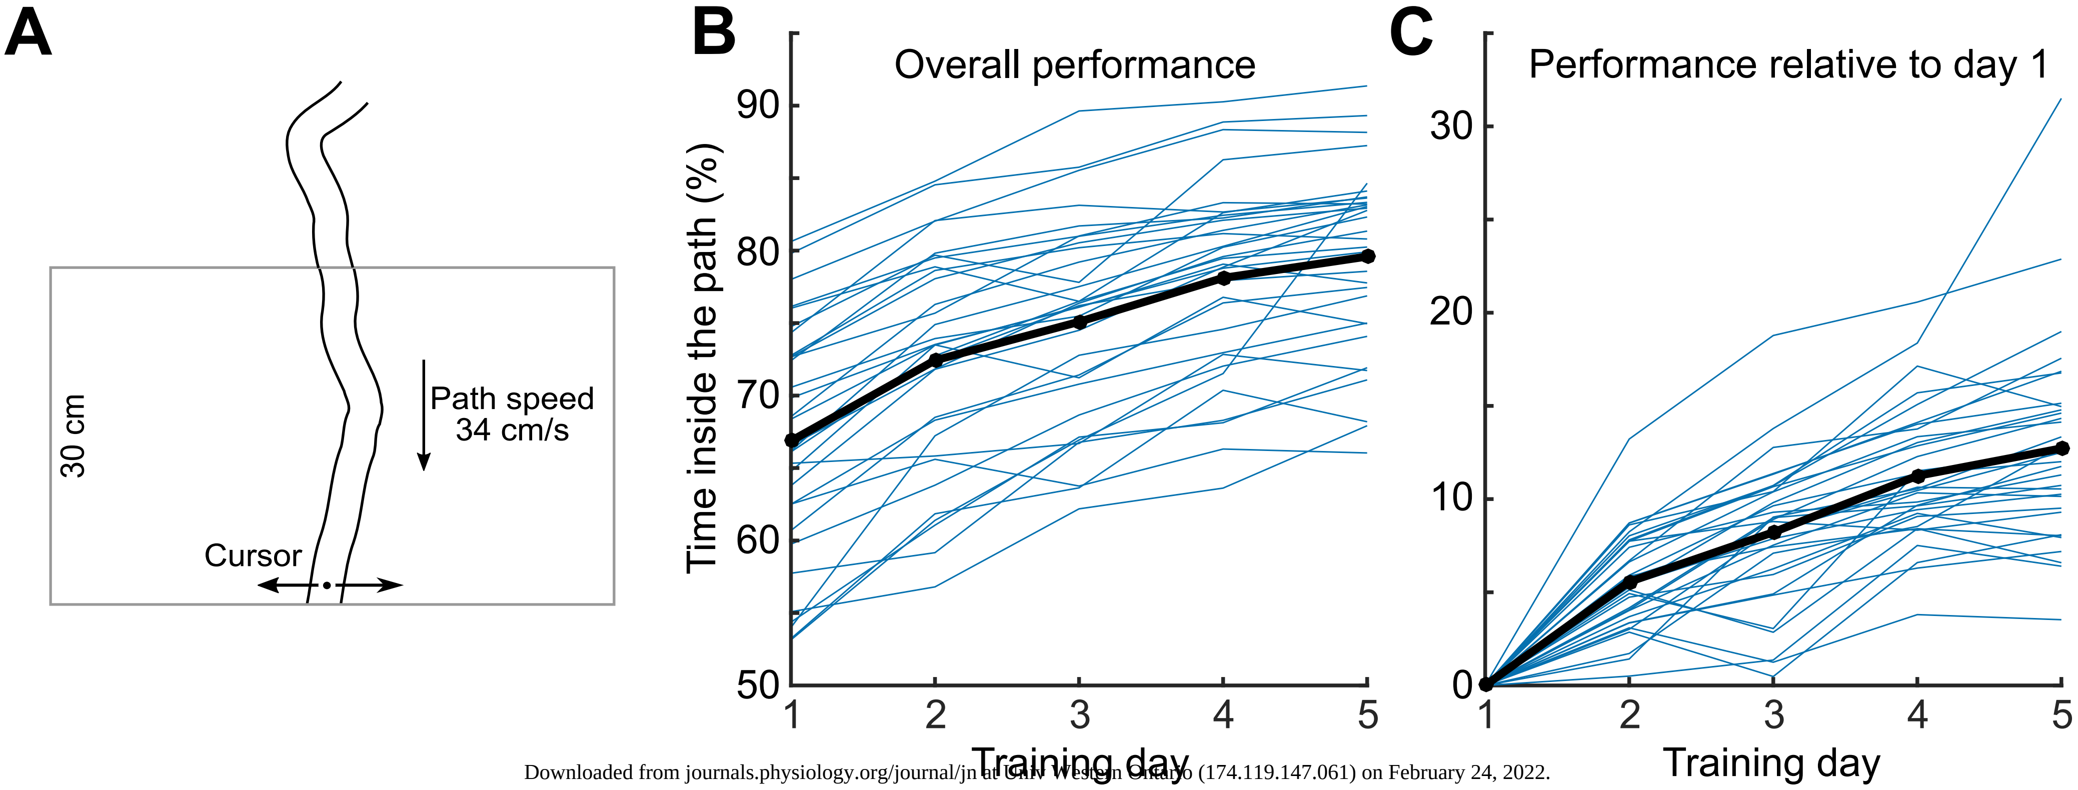
A
30 cm
Path speed
34 cm/s
Cursor
B
Overall performance
Time inside the path (%)
90
80
70
60
50
1
2
3
4
5
Training day
C
Performance relative to day 1
30
20
10
0
1
2
3
4
5
Training day

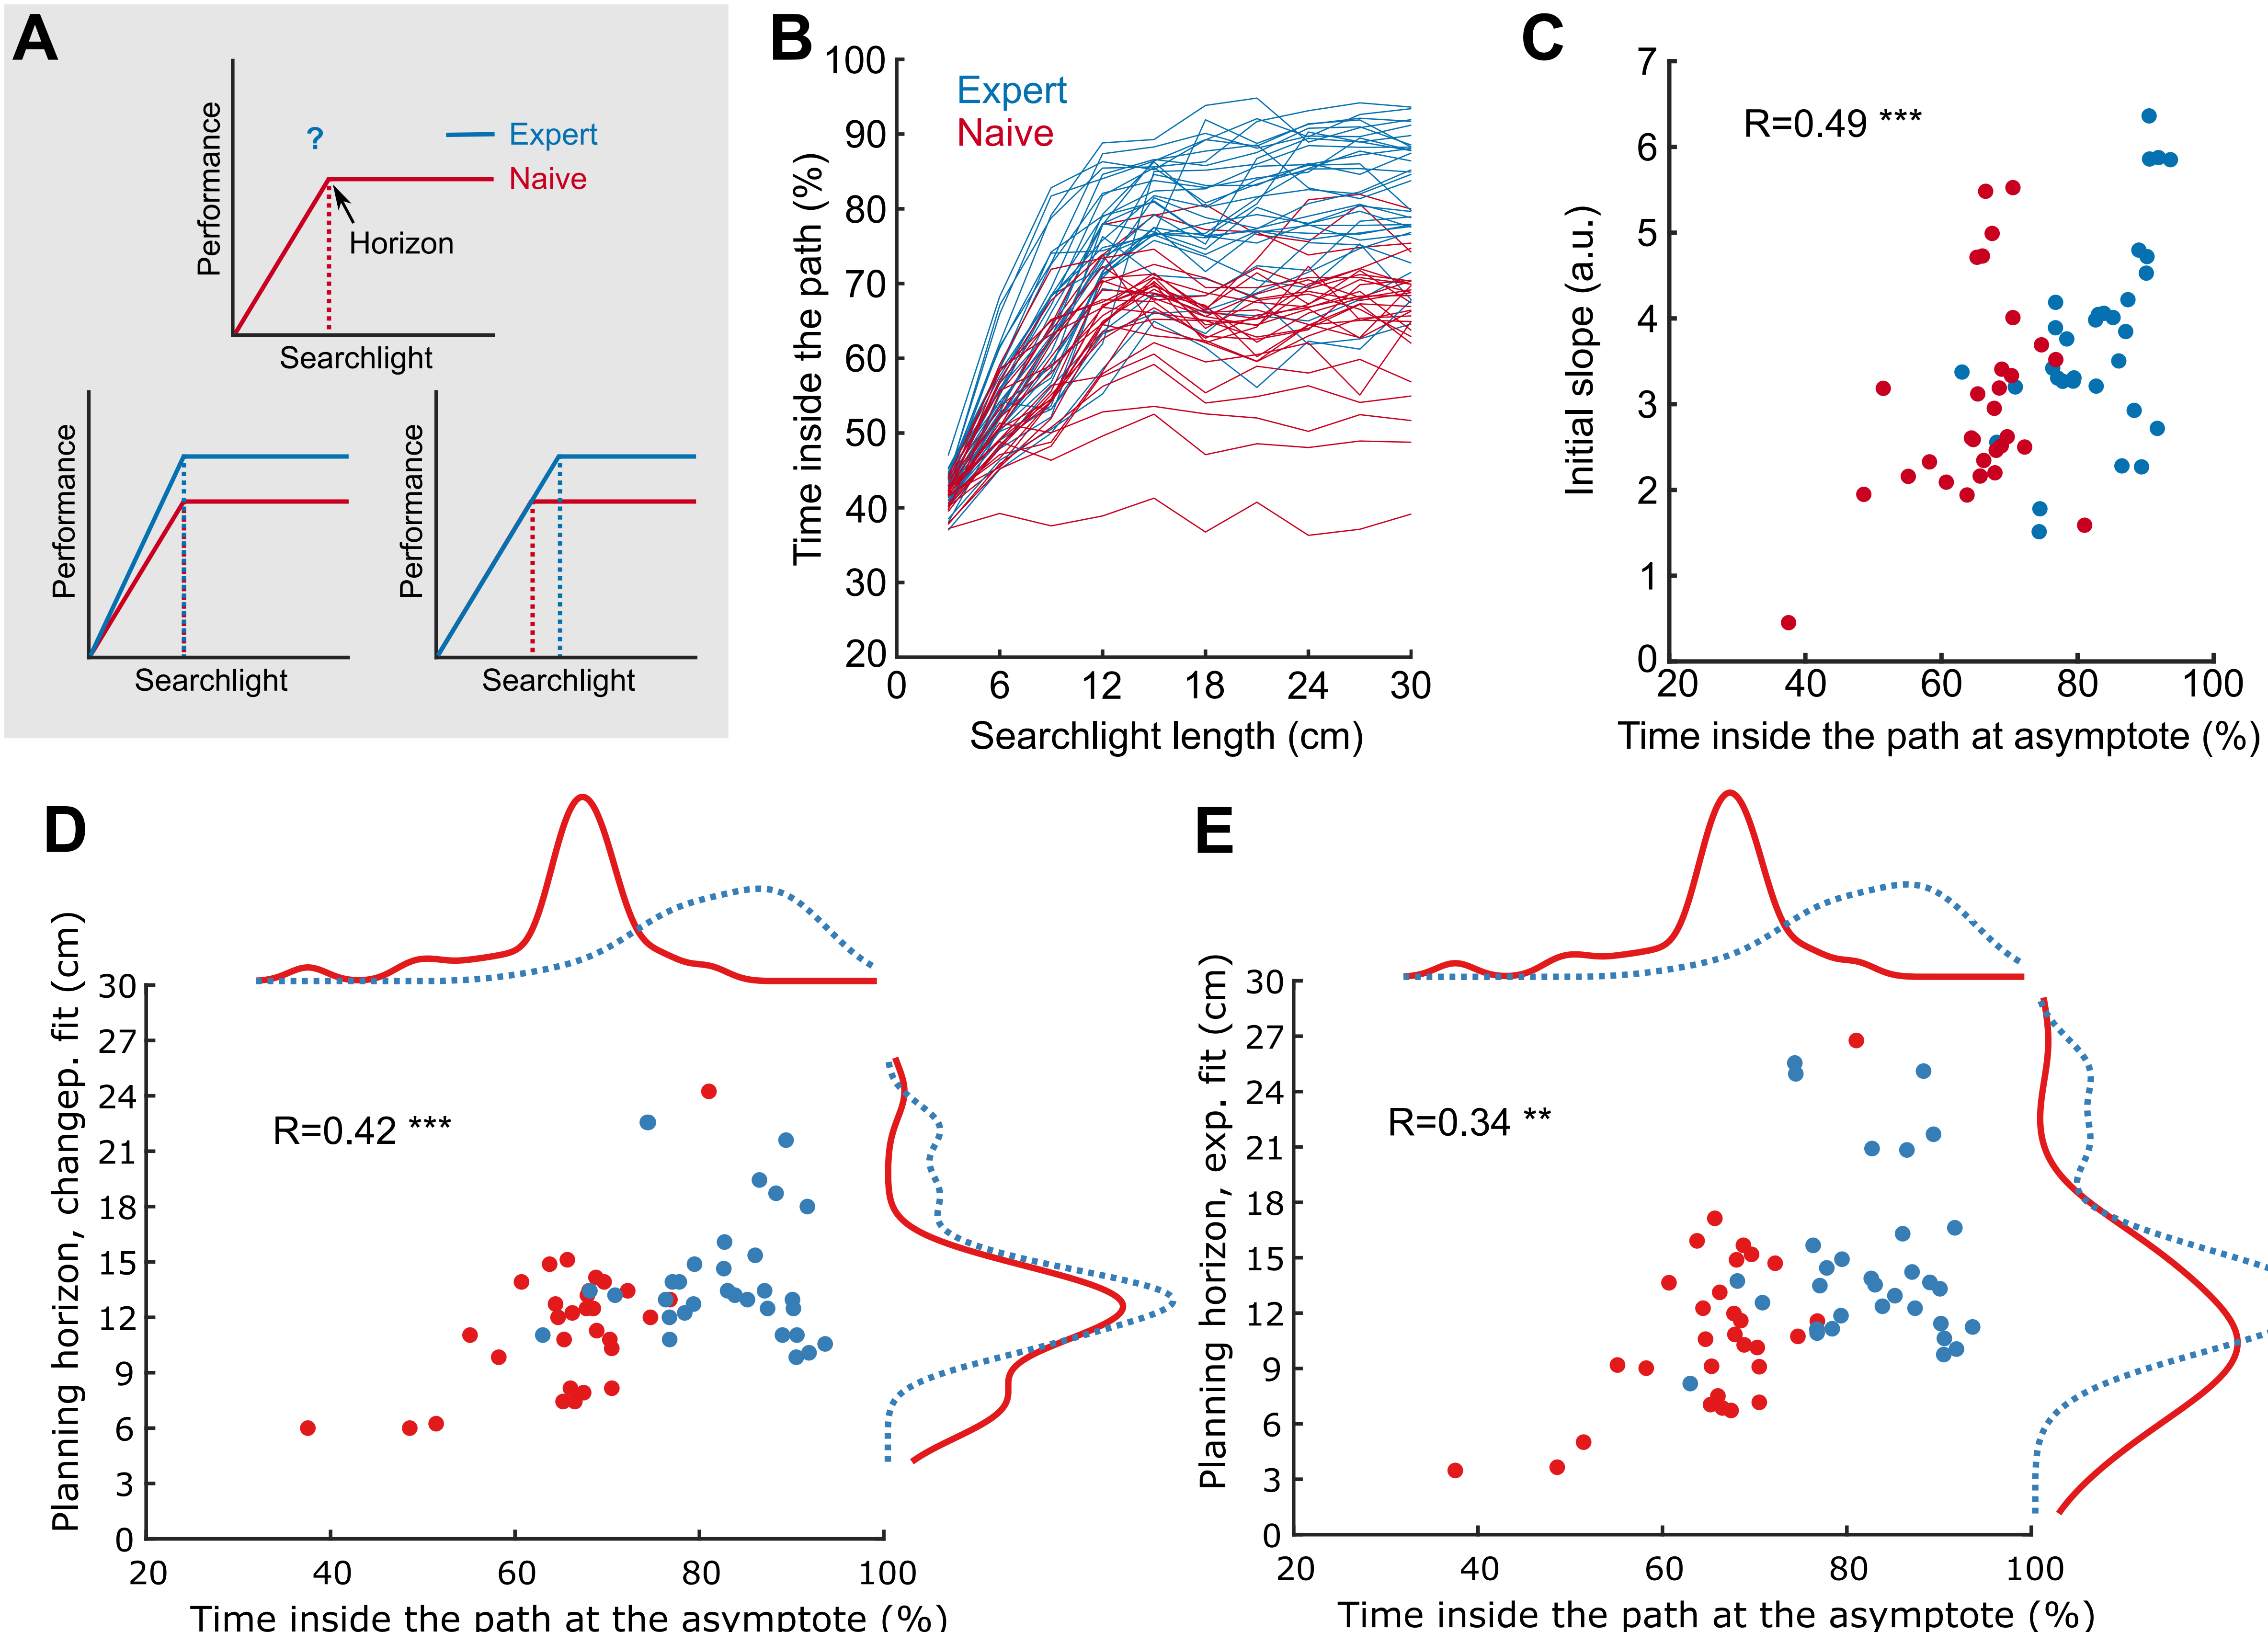

A
Performance
?
Expert
Naive
Horizon
Searchlight
Performance
Searchlight
Performance
Searchlight
B
Expert
Naive
Time inside the path (%)
Searchlight length (cm)
C
R=0.49 ***
Initial slope (a.u.)
Time inside the path at asymptote (%)
D
R=0.42 ***
Planning horizon, changep. fit (cm)
Time inside the path at the asymptote (%)
E
R=0.34 **
Planning horizon, exp. fit (cm)
Time inside the path at the asymptote (%)

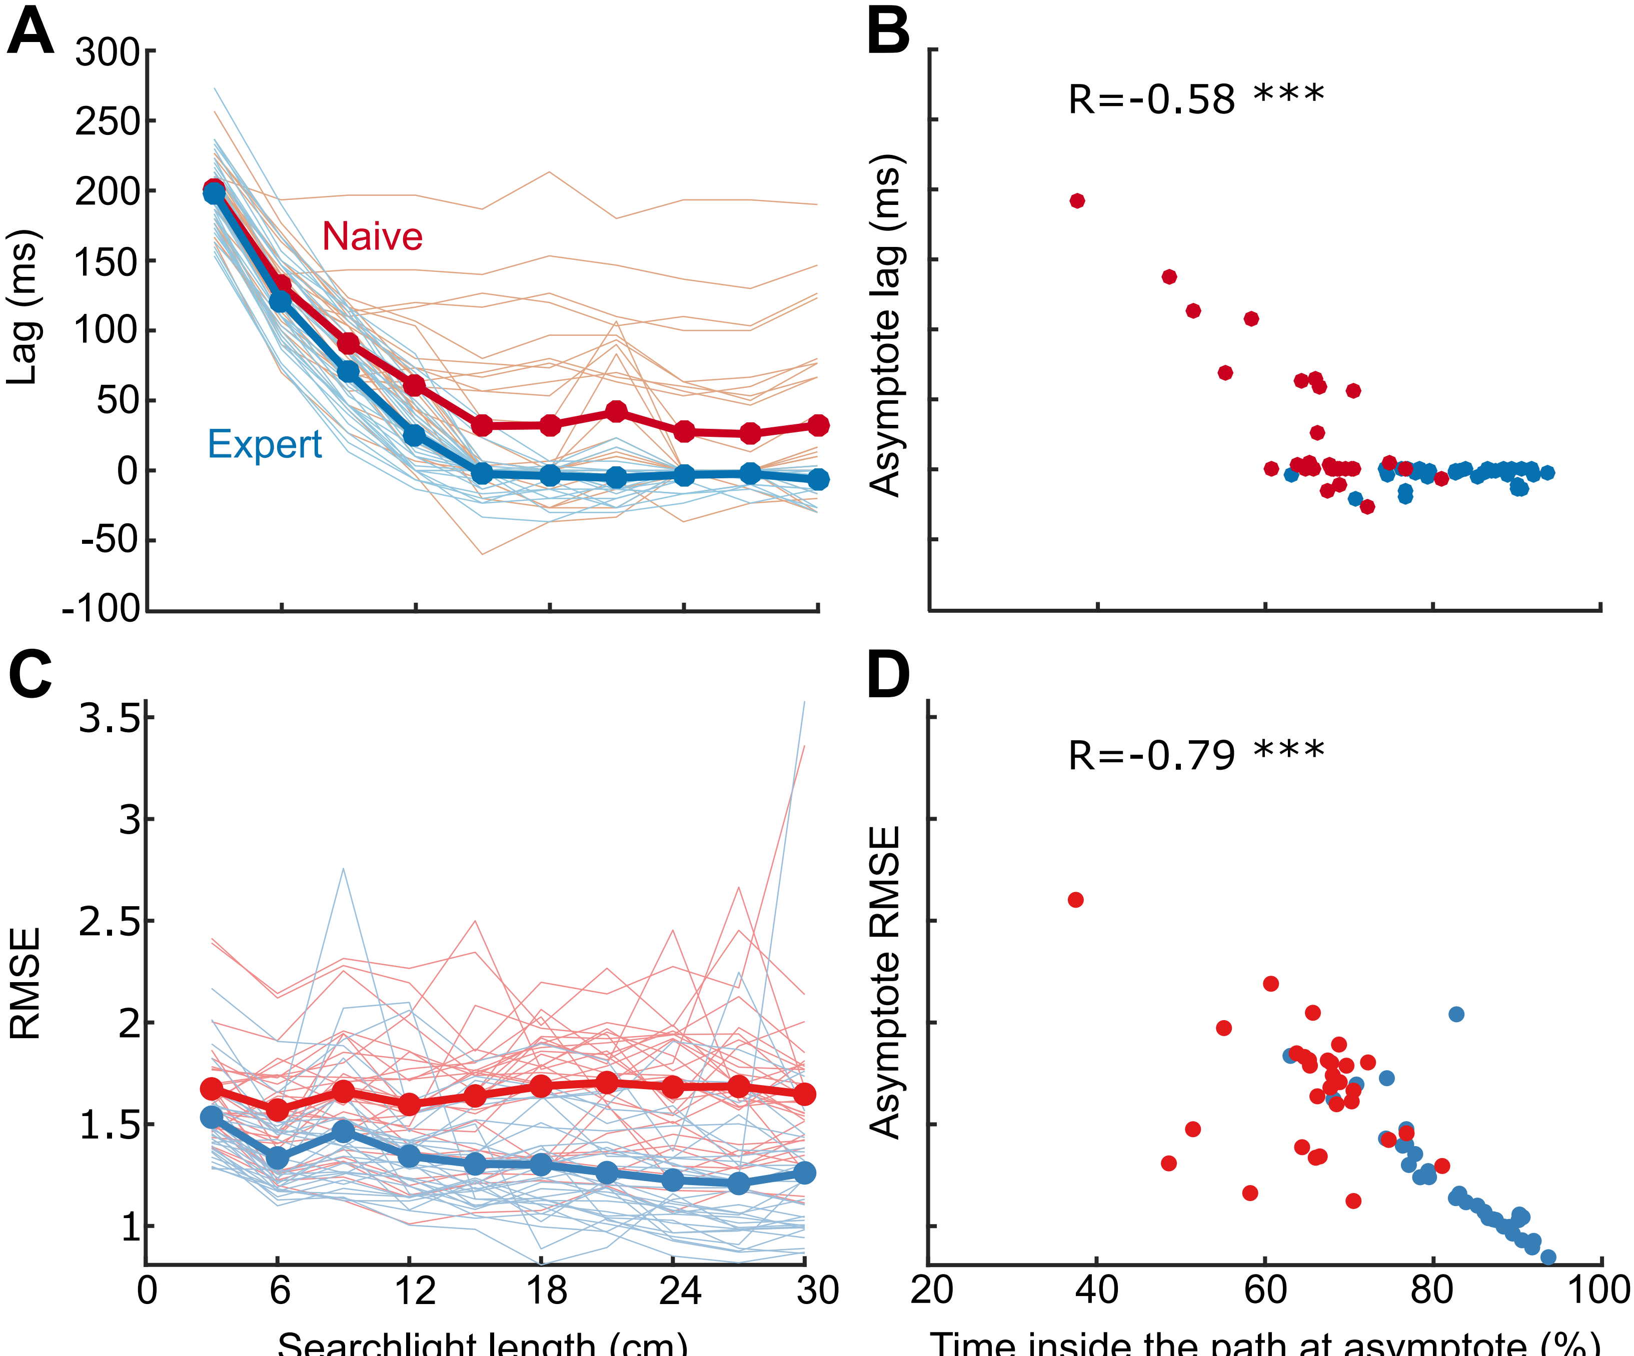
A
B
C
D
Naive
Expert
R=-0.58 ***
R=-0.79 ***
Lag (ms)
Asymptote lag (ms)
RMSE
Asymptote RMSE
Searchlight length (cm)
Time inside the path at asymptote (%)
300
250
200
150
100
50
0
-50
-100
3.5
3
2.5
2
1.5
1
0
6
12
18
24
30
20
40
60
80
100

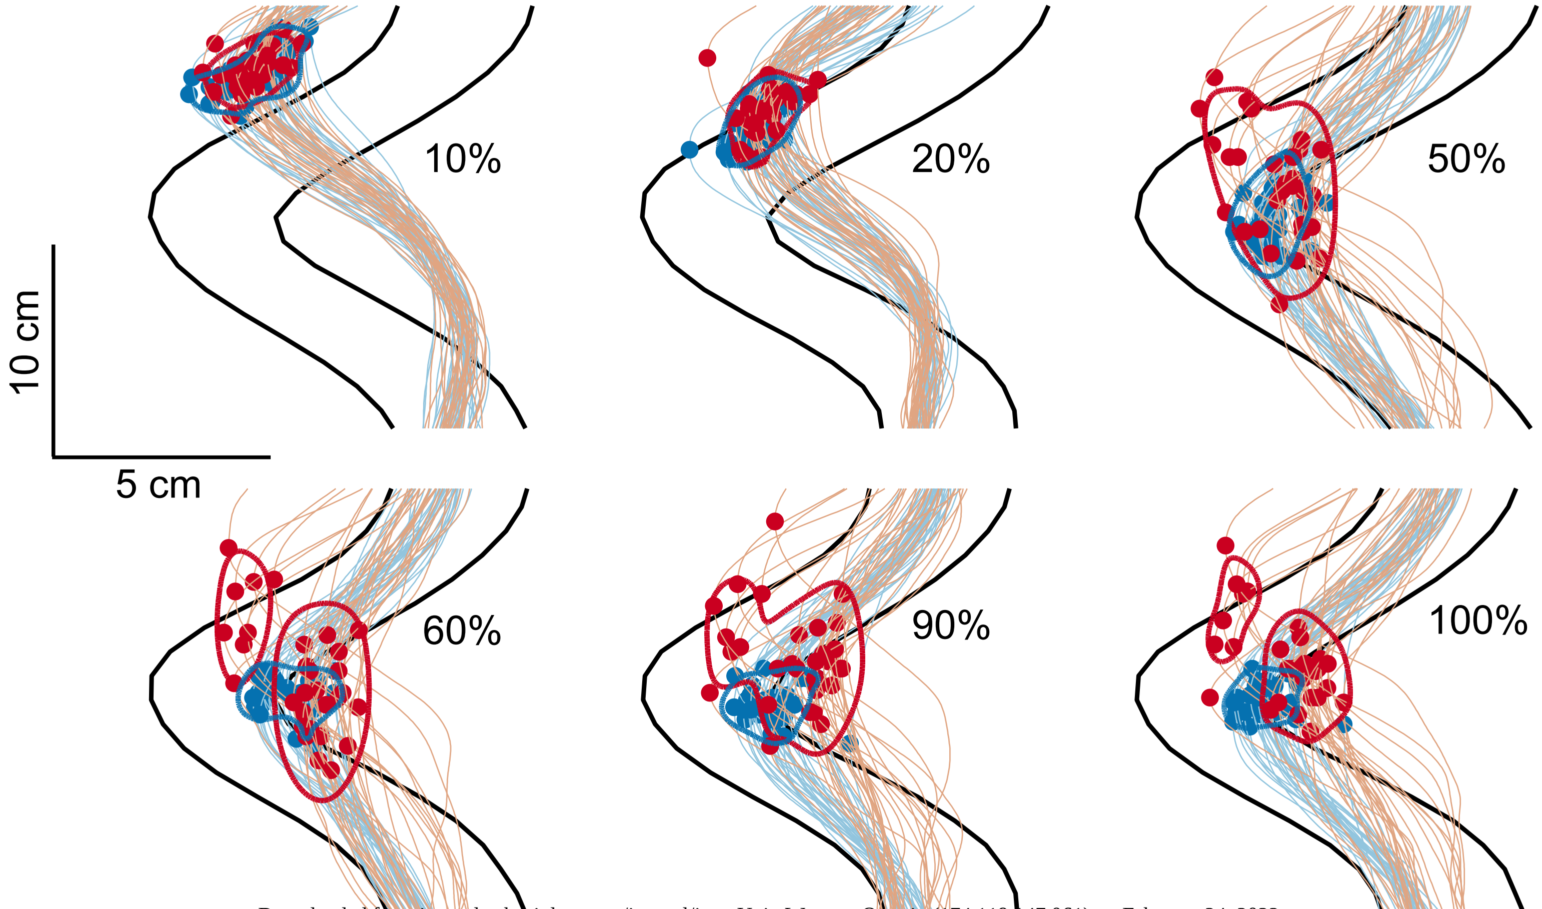
10%
20%
50%
60%
90%
100%
10 cm
5 cm

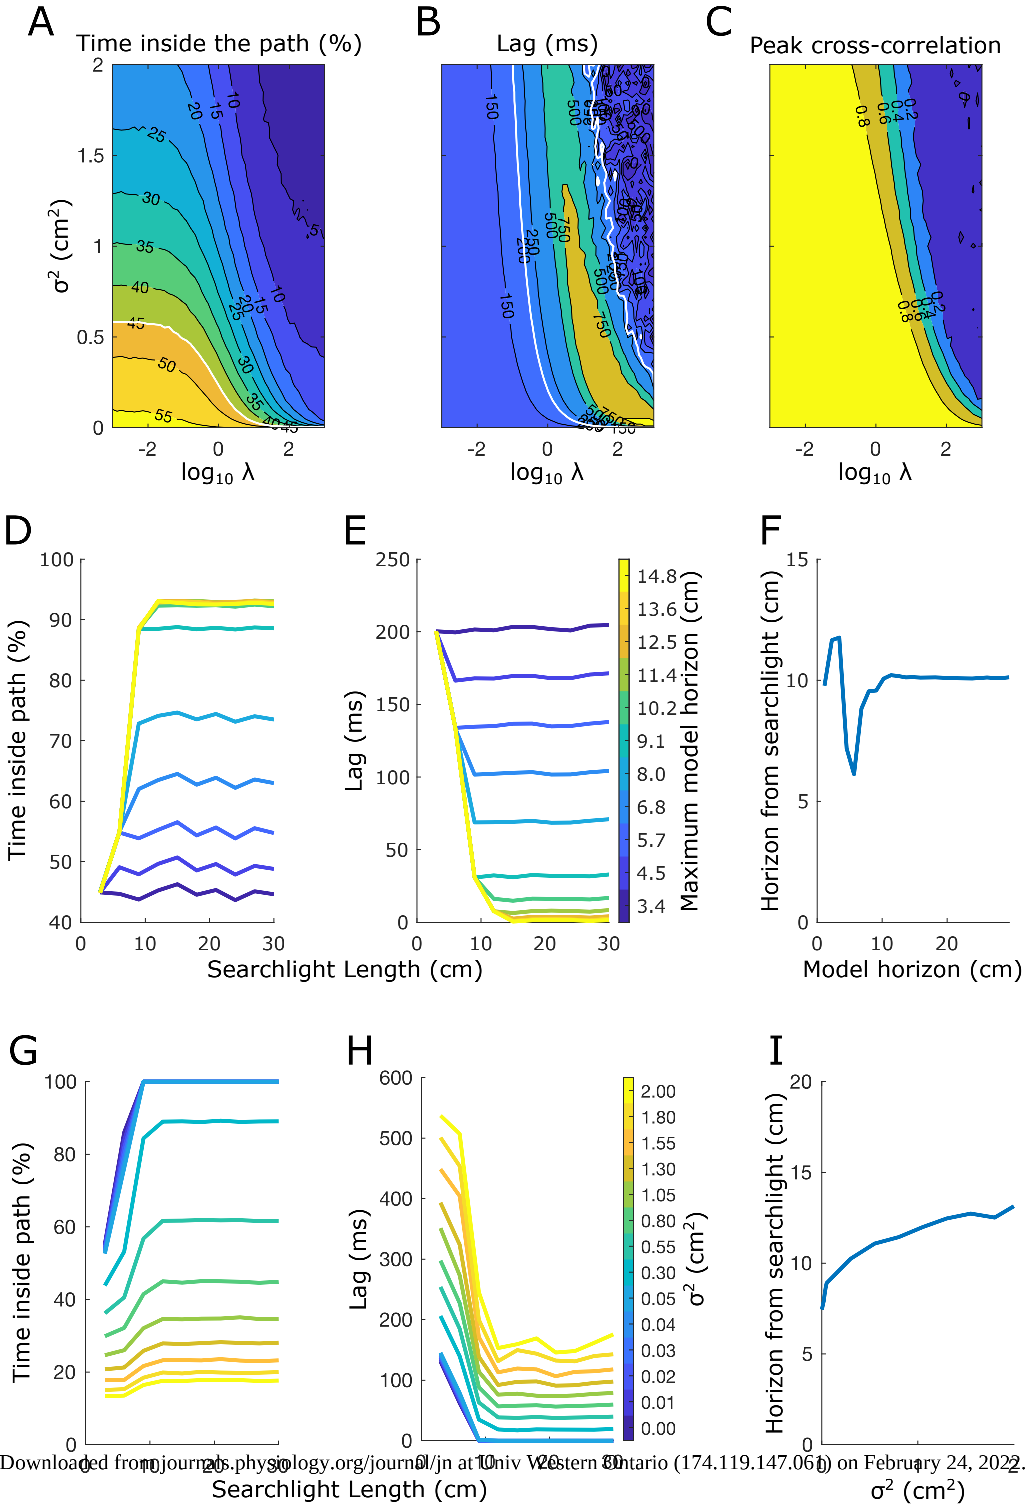
A
Time inside the path (%)
σ² (cm²)
log₁₀ λ
B
Lag (ms)
log₁₀ λ
C
Peak cross-correlation
log₁₀ λ
D
Time inside path (%)
E
Lag (ms)
Maximum model horizon (cm)
14.8
13.6
12.5
11.4
10.2
9.1
8.0
6.8
5.7
4.5
3.4
Searchlight Length (cm)
F
Horizon from searchlight (cm)
Model horizon (cm)
G
Time inside path (%)
H
Lag (ms)
σ² (cm²)
2.00
1.80
1.55
1.30
1.05
0.80
0.55
0.30
0.05
0.04
0.03
0.02
0.01
0.00
Searchlight Length (cm)
I
Horizon from searchlight (cm)
σ² (cm²)

Motor skill learning decreases movement variability and increases planning horizon

Methods

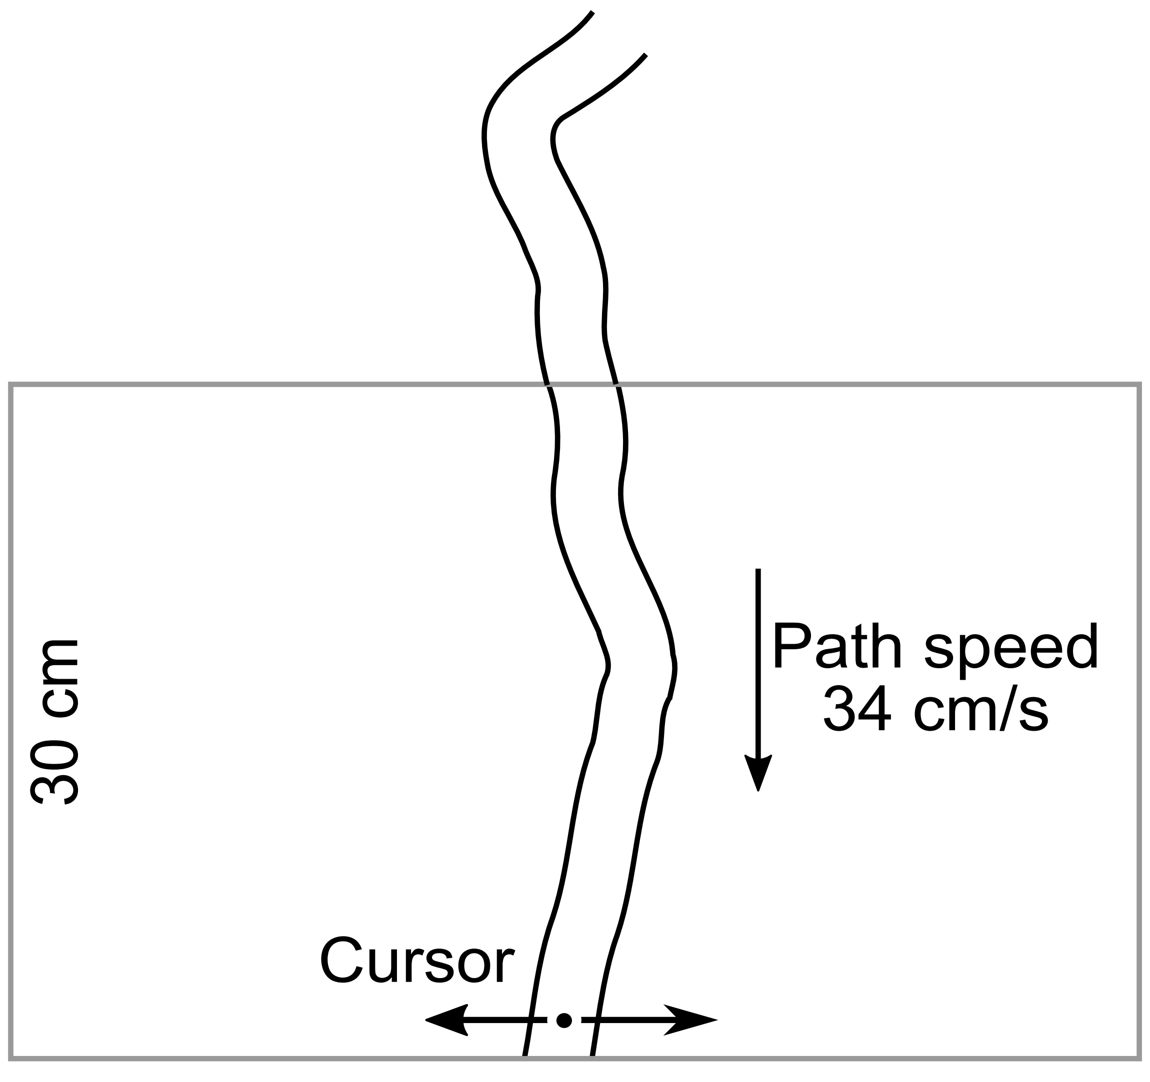

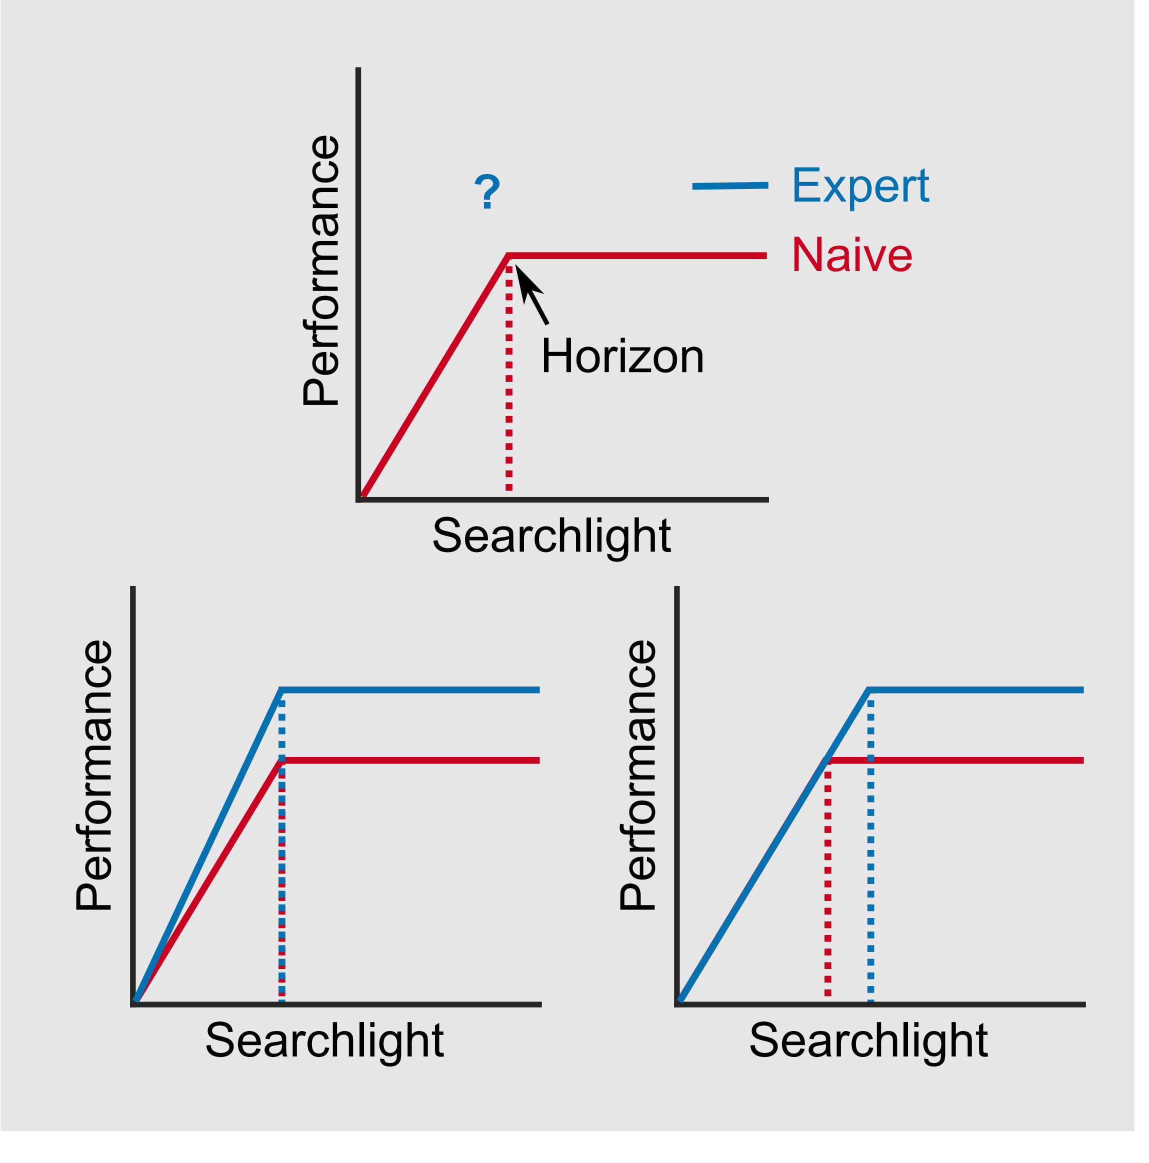

Outcomes

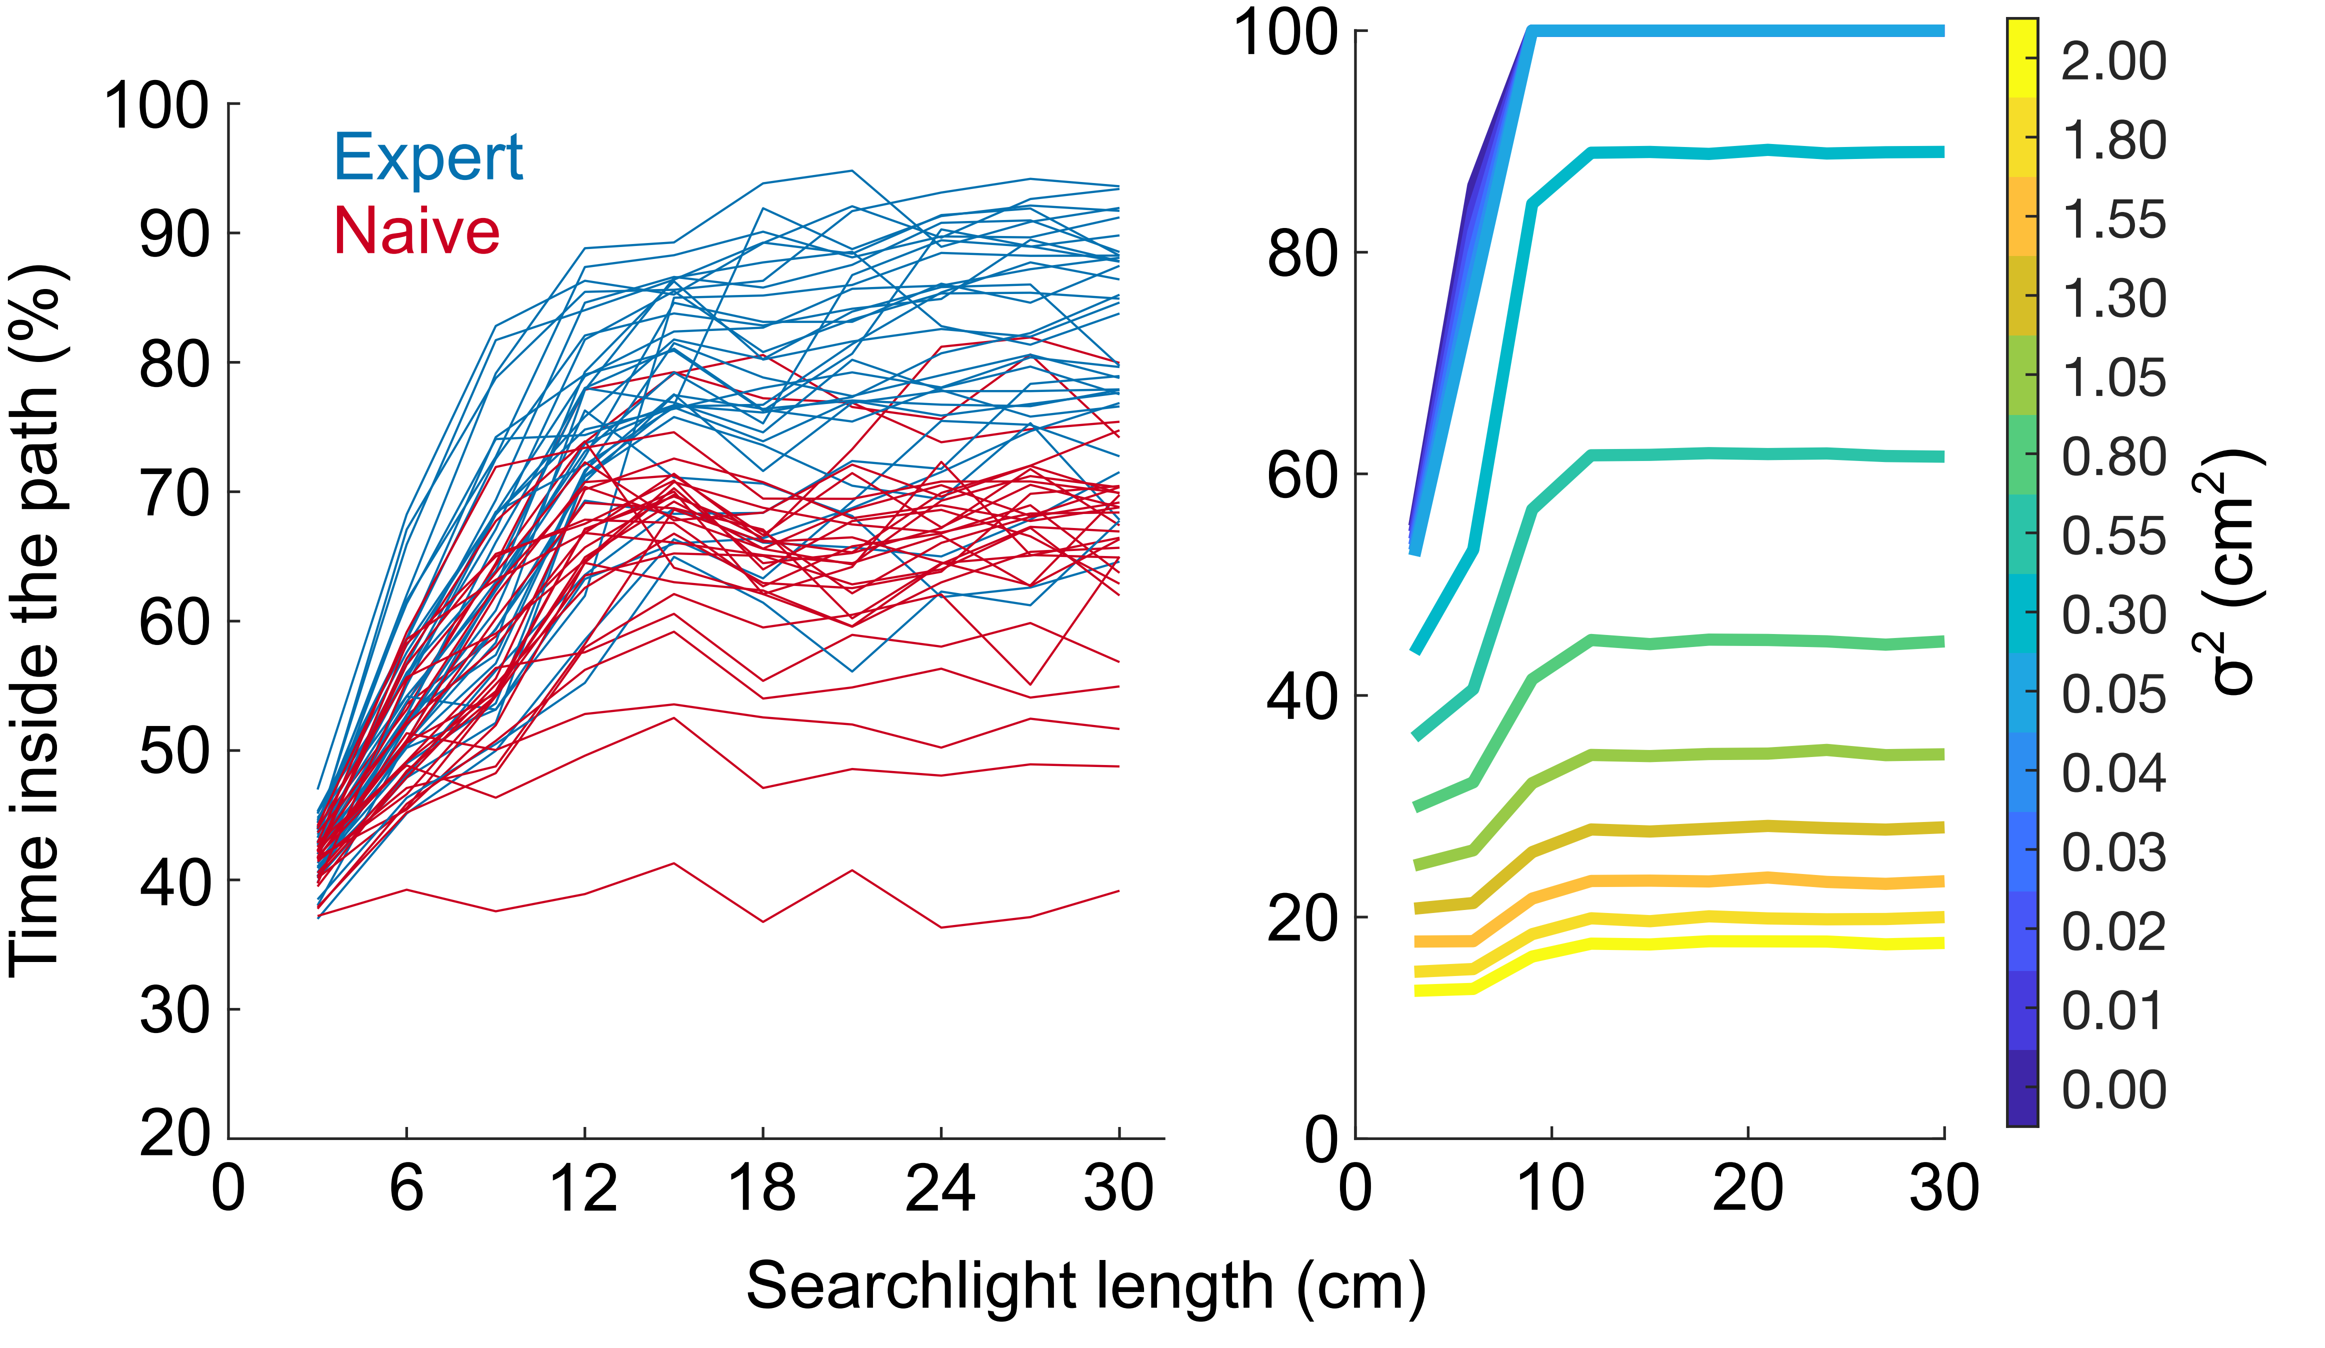

Conclusion

Experts had a lower movement variability and took a longer section of the future path into account when performing the task. A receding horizon control model that increases with tracking still quantitatively captured the movement behavior.