



Published in final edited form as:

Exp Brain Res. 2009 February ; 193(1): 69–83. doi:10.1007/s00221-008-1596-1.

Variability In Motor Learning: Relocating, Channeling and Reducing Noise

R.G. Cohen¹ and D. Sternad²

¹ Department of Neurology, Oregon Health Sciences University

² Departments of Biology, Electrical & Computer Engineering, and Physics, Northeastern University

Abstract

Variability in motor performance decreases with practice but is never entirely eliminated, due in part to inherent motor noise. The present study develops a method that quantifies how performers can shape their performance to minimize the effects of motor noise on the result of the movement.

Adopting a statistical approach on sets of data, the method quantifies three components of variability (Tolerance, Noise, and Covariation) as costs with respect to optimal performance. *T-Cost* quantifies how much the result could be improved if the location of the data were optimal, *N-Cost* compares actual results to results with optimal dispersion at the same location, and *C-Cost* represents how much improvement stands to be gained if the data covaried optimally. The *TNC-Cost* analysis is applied to examine the learning of a throwing task that participants practiced for 6 or 15 days. Using a virtual set-up, 15 participants threw a pendular projectile in a simulated concentric force field to hit a target. Two variables, angle and velocity at release, fully determined the projectile's trajectory and thereby the accuracy of the throw. The task is redundant and the successful solutions define a nonlinear manifold. Analysis of experimental results indicated that all three components were present and that all three decreased across practice. Changes in *T-Cost* were considerable at the beginning of practice; *C-Cost* and *N-Cost* diminished more slowly, with *N-Cost* remaining the highest. These results showed that performance variability can be reduced by three routes: by tuning tolerance, covariation and noise in execution. We speculate that by exploiting *T-Cost* and *C-Cost*, participants minimize the effects of inevitable intrinsic noise.

Keywords

motor learning; variability; noise; sensitivity; skill acquisition

It is widely acknowledged that improvement in skilled performance with practice includes a decrease in variability. However, even in expert performance, variability is never entirely eliminated. One explanation for the ubiquitous presence of variability is that the nervous system is inherently noisy. The sensorimotor system is a complex dynamical system with processes

Correspondence Address: Dagmar Sternad, Departments of Biology, Electrical & Computer Engineering, and Physics, 134 Mugar Life Science Building, Northeastern University, 360 Huntington Avenue, Boston MA 02115, Phone: 617-373-5093, email: dagmar@neu.edu.

²A manifold is a mathematical term from differential geometry generally referring to a collection or set of objects (James & James, 1992, 5th ed.). The set of solutions in the skittles task is a smooth differentiable manifold whether or not the execution variables have different units. The interested reader can find definitions at any level of detail in relevant mathematical texts. For example, James and James (1992, 5th ed): "A topological manifold of dimension n is a topological space such that each point has a neighborhood which is homeomorphic to the interior of a sphere in Euclidean space of dimension n ." The term has received many qualifications, including connected or disconnected, compact or noncompact, etc.

³The model with the two orthogonal springs is equivalent to one with a single spring rotating around a center pivot. However, both models are approximations as they do not include vertical elevation as is present in the real skittles task.

at many different time scales, in which fluctuations percolate across levels and are manifested in overt behavior as “noise” or rather variability. Sources of such variability have been identified in the basal ganglia circuits, premotor cortex, and in the recruitment properties of motor units, both in the planning and the execution of movements (Churchland, Afshar, & Shenoy, 2006; Jones, Hamilton, & Wolpert, 2001; Ölveczky, Andalman, & Fee, 2005; van Beers, Haggard, & Wolpert, 2004). Given precision requirements in many everyday tasks, much attention has been given to the problem of how the sensorimotor system minimizes unwanted effects of its own intrinsic noise on performance (Fitts, 1954; Harris & Wolpert, 1998; Schmidt, Zelaznik, Hawkins, Frank, & Quinn, 1979).

A more positive interpretation for the presence of such “noise” is that it affords exploration of the task and therefore aids in finding the best solutions to achieve the desired result (Newell & McDonald, 1992; Ölveczky et al., 2005; Riccio, 1993). Simulation studies testing learning algorithms in artificial neural networks have demonstrated the positive effects of noise in accelerating the learning rate and avoiding local minima, i.e., suboptimal solutions (Burton & Mpitsos, 1992; Geman, Bienenstock, & Doursat, 1992; Metropolis, Rosenbluth, Rosenbluth, Teller, & Teller, 1953). Another positive perspective on variability is that even when the system is at a skilled level, a given amount of noise affords flexibility and maneuverability (Beek & van Santvoord, 1992; Full, Kubow, Schmitt, Holmes, & Koditschek, 2002; Hasan, 2005). Given these contrasting perspectives, there is a need for more sophisticated methods for evaluating variability in skilled behavior and differentiating between positive and negative effects of variability on performance. The aim of this project is to develop a novel method to quantify different components of variability and apply it to experimental data acquired during the acquisition of a skill that requires precision.

Almost any motor task in the real world contains more degrees of freedom than are strictly required, such that a desired result can be achieved in more than one way. This insight is generally attributed to Bernstein, who, observing the hammering of blacksmiths, pointed out that multi-joint movements were executed slightly differently at every repetition, while the end points of the hammer strokes were remarkably consistent (Bernstein (1967, 1996); see also Lashley (1930)). His observations stimulated studies demonstrating covariation amongst redundant degrees of freedom that made performance less variable over repetitions. To understand why, consider a simplified example. If a task is defined such that X and Y must sum to 5 in order to achieve perfect performance, and both X and Y vary between 0 and 5 with means of 2.5, then a negative linear relationship between X and Y increases the likelihood of achieving a perfect result. In the context of a motor task, X and Y could represent forces exerted by different muscles or limbs (Latash, Scholz, Danion, & Schöner, 2001), or angles traveled by different joints (Scholz, Schöner, & Latash, 2000), or angles and velocities of release of a ball (Müller & Sternad, 2004a). However, bear in mind that the relation between X and Y may be much more complex than the simple negative linear correlation in the above example. Note that we use the term *covariation* for any correlated variation between two or more variables (Merriam Webster) while *covariance* is reserved for the well-defined statistical concept.

Recognizing the importance of motor equivalence, several research groups have developed methods to describe how functional covariation among elements reduces the negative effects of variability on the result. For example, Kudo and colleagues (2000) have employed a surrogate method to assess the extent of covariation between variables, similar to Müller and Sternad (2003). The UCM approach (Scholz et al., 2000) employs null space analysis, in adaptation of mathematical approaches developed for the control of multi-degree-of-freedom systems in robotics (Craig, 1986; Liégeois, 1977). This approach examines covariance in execution, most frequently amongst joint angles, and does not directly relate this variability to accuracy in a task-defined variable. Very few approaches have related execution to results

(Cusumano & Cesari, 2006; Kudo, Tsutsui, Ishikura, Ito, & Yamamoto, 2000; Martin, Gregor, Norris, & Thach, 2001; Müller & Sternad, 2003, 2004b; Stimpel, 1933).

The present approach is based on a two-level description of the task; it presents a framework for examining the relation between variability in *execution* and variability in *result*¹. Note that in the blacksmith example above, the execution (the trajectories of the joint angles and the hammer) may be considered separately from the result (the contact location of the hammer with respect to the target location), but the execution necessarily determines the result. Similarly, in a throwing task, the position and velocity of the arm at the moment of ball release are the *execution variables* that fully determine the trajectory of the ball and thereby the result, i.e., throwing accuracy or distance. The *result variable* can be defined as any outcome based on the trajectory, such as the maximum height reached by the ball, the total distance it travels, or how close it comes to a target. This two-level description is particularly interesting when, as in the examples offered above, the dimensions of the execution space (defined by the number of execution variables) are greater in number than the dimensions of the selected result variables.

In addition to the exploitation of the null space by covariation between the execution variables, there are other ways that variability in the execution of a task can benefit performance. For instance, variability will necessarily arise at the beginning of a learning process when an actor explores different strategies before finding the best solutions (Müller & Sternad, 2003, Müller & Sternad, 2004). Despite the general recognition of this important stage in learning and development, remarkably little has been done to describe this process quantitatively. This is partly because it has proven difficult to extract any systematicity in this exploration process and consequently quantify it (Newell & McDonald, 1992). A similar related problem is the observation that biological systems may change their movements in order to avoid placing uniform and persistent stress on joints and muscles (Duarte & Zatsiorsky, 1999). Hence, random fluctuations or noise may also have benefits to performance. Given all these potentially functional aspects of variability, a fine-grained decomposition of performance into potentially separate components is an interesting and as-of-yet insufficiently addressed problem.

The present study builds on the insights of Müller and Sternad (2003, Müller and Sternad (2004), who proposed a decomposition method to quantify three components of variability: tolerance, noise, and covariation (*TNC*). The three components express three conceptually different routes to how performance can be improved. While the previous method quantified *T*, *N*, and *C* in terms of relative changes across practice sessions, the new approach takes an optimization perspective that formulates the three components as costs with respect to ideal performance. A new numerical method is developed that assesses to what degree the result of non-optimal performance could be improved if one of these components were optimized. Details of the new method are presented after the introduction of the specific task and the concepts developed for representing task performance. A direct comparison of the two methods is included in the discussion.

The task is a simplified, virtual version of the British pub game called “Skittles,” previously introduced by Müller and Sternad (2004). In the pub version of the game, a ball hangs from a pole, and one or more targets (the skittles) are placed on the far side of the pole from the player. The virtual version of the game emulates the three-dimensional dynamics of the pub game by providing the player with a top-down view of the workspace, as shown in Figure 1A. The player’s forearm rests on a manipulandum that is fixed at the elbow, limiting movement to a

¹The words “result” and “result variable” are used as a general term to designate any measure that quantifies how satisfactorily a task was achieved, such as error, score, difference from target, or even a variability estimate if sets of data were evaluated for their result. In the present study we only consider one such measure: distance to target.

horizontal rotation around the elbow joint. The player attempts to throw the ball so that it swings around the pole and hits a single target. The player controls the timing of the ball release by extending a finger to open a contact switch. The task is redundant because two execution variables (angle and velocity of manipulandum at release) co-determine a single result variable. The result variable, distance to target, is defined as the minimum distance between the trajectory and the center of the target skittle. Trial 1 in the figure illustrates how the distance to the target is defined. Trials 2 and 3 show how different release angles and velocities produce different ball trajectories that both achieve a perfect hit.

Figure 1B represents the same three throws in a space spanned by the two execution variables, release angle and release velocity. The result attained by each pair of execution variables is represented by the gray shades, and the white U-shaped area represents the subset of release angles and velocities that lead to a direct hit with error less than 1.2 cm. The subset of perfect solutions will be called the *solution manifold*. Note that trials 2 and 3, whose trajectories lead to hits with near-zero error, are on the solution manifold, while trial 1 is in a grey area, indicating a larger error. Characteristics of the result space, including the size, shape, and location of the solution manifold, are determined by factors in the *work space*, namely the location and size of the target and the center post. For examples of other result spaces see (Hu, Müller, & Sternad, submitted; Müller & Sternad, 2004b).

The *TNC*-approach differs from other approaches to variability in several respects. One feature central to the analysis of performance presented here is the focus on sets of trials rather than the time series of trials. Each set of trials has a location, a size, and a shape in the execution space. Second, the approach permits evaluation of tolerance (location), covariation, and noise (dispersion) of a set of data, and changes in these properties give insight into the changes due to practice. Third, a distinguishing feature of the *TNC*-approach is that the sensitivity of different regions of the *solution manifold* to error is explicitly taken into account. In Figure 1B, regions of the solution manifold with greater sensitivity or tolerance to error are surrounded by larger areas in light gray shades. Such regions are less prone to error and are referred to as having higher *tolerance* to error. Hence, aiming at more tolerant regions of the execution space reduces the effect of variability in execution on the result.

There were two goals to the study reported here. The first was to develop the *TNC-Cost* analysis. The three conceptually distinct components that are proposed as contributors to learning are reconceived here as costs to performance for a given data set. Tolerance, or the deviation from the most tolerant region of execution/result space, is defined as *T-Cost*, which is the cost to performance of not having found that region – that is, of not being at the best location. As will be clearer from the detailed description below, this is a new non-local concept not considered in other approaches. Covariation, quantifying the ability of participants to exploit the redundancy inherent to the execution/result space, is defined as *C-Cost*, which is the cost to performance of not exploiting that redundancy – that is, a failure to use optimal covariation. Noise reduction, or the reduction of dispersion in the execution variables, is defined as *N-Cost*, which is the cost to performance due to noise in the execution variables. Note that *N-Cost* is not identical to noise in the execution variables alone; rather, noise is evaluated by mapping its effect into result space. Optimal noise is defined by the associated optimal result. For all three costs, their quantification relies on the creation of virtual data sets that are optimized in terms of one component while the others are held constant.

The second goal of this study was to examine the relative importance of and relations among different costs to performance during skill acquisition, in both a group of healthy participants with average expertise and three expert throwers. The inclusion of three highly skilled participants that practiced for 15 days allowed us to scrutinize later stages of skill improvement.

Method

Participants

A total of 12 right-handed participants (2 female, 10 male) were tested, after having given informed consent to participate in the experiment. All were members of the Pennsylvania State University community. They ranged in age from 21 to 48 years. The 5 undergraduate students received course credit for their participation. The 8 graduate students, one faculty and one staff member were not compensated for their time. Three participants (referred to hereafter as the “expert” participants) had more than average expertise: One was an avid cricket player with extensive throwing experience, one was a member of the university’s Ultimate Frisbee team, and one was all-around exceptional athlete. The protocol was approved by the Institutional Review Board of the Pennsylvania State University.

Task and Apparatus

Emulating the real game of skittles, in the virtual experimental task the participant attempted to throw a ball such that it swung around a center pole to hit a single target located at the far side. The participant stood approximately 60 cm from a projection screen (width: 2.50 m, height: 1.80 m) as seen in Figure 2. A computer-generated display of a top-down view of the work space was projected onto the vertical screen (as illustrated in Figure 1a). The post in the center of the work space was represented by a circle of 16 cm diameter, whose center was located 1.73 m above the floor. A circular target of 3 cm diameter was presented, with its center 20 cm to the right and 50 cm above the center of the post. The virtual arm was represented as a solid bar of 12 cm length, fixed at one end. The fixed end (rotation point) of the virtual arm was centered 50 cm below the center of the post, and a circle of 3 cm diameter representing the ball was attached to the free end of the virtual arm.

Standing in front of the screen, the participant rested his or her forearm on the manipulandum that was fixed to a vertical support, which was adjusted to a comfortable height for each participant. The horizontal manipulandum pivoted around an axle centered directly underneath the elbow joint. The metal support was padded with foam and the elbow was fixed with a Velcro strap. Rotations of the arm were measured by a 3-turn potentiometer with a sampling rate of 700 Hz (the maximum allowed by the online computations for the display). At the free end of the metal arm, a tennis ball was attached with a pressure-sensitive switch glued to its surface. The switch was located on the ball such that the index finger covered the switch in a natural grasp configuration. The participant grasped the ball with his or her right hand, closed the switch with the index finger, then swung the arm in an outward horizontal circular motion. Releasing the switch contact triggered the release of the virtual ball, which traversed a trajectory initialized by the angle and velocity of the participant’s arm at the moment of release. Both the movements of the arm and the simulated trajectory of the ball were displayed on the screen in real time. The ball’s trajectory, as determined by the simulated physics of the task, described an elliptic path around the pole. This trajectory was not immediately intuitive to participants, and they had to learn the mapping between the real arm movements and the ball’s trajectories in the projected work space. Hence, the task was novel even for the participant with extensive throwing experience.

Instructions and Procedure

Participants stood with their right shoulders facing the display. The experimenter instructed participants to throw the ball in a clockwise direction toward the screen as if performing a Frisbee backhand, and showed them approximately where to release the ball. Although the throwing action could be performed in several other ways, participants were not permitted to explore other strategies. After each throw, an enlarged image of a part of the trajectory and the target was shown to provide feedback to the participant about the accuracy of the throw. No

quantitative feedback was given. However, if the trajectory passed within 1.3 cm of the center of the target, the target color changed from yellow to red. Participants were encouraged to achieve as many of those “direct” hits as possible.

Participants performed 180 throws per day. After each 60 throws, participants were given a short break. Each participant performed experimental sessions on six days, separated by one or two days on average. Three expert participants volunteered to practice for 15 days, which allowed us to examine changes in the structure of variability during the fine-tuning phase of performance.

Data Analysis and Dependent Measures

The ball trajectories were simulated online based on the measured angle and derived velocity at the moment of release. To reduce contamination from measurement noise in the picking of release angle and velocity we first fitted the last 10 samples of the angular position before the moment of release with a straight line. (Given that we limited the fits to samples before the release moment, cubic splines did not improve the fit.) This linear regression was used to pick angle and velocity at the release moment. To evaluate the result, the minimum distance between the trajectory and the center of the target was calculated. The elliptic trajectories of the ball were generated by a two-dimensional model in which the ball was attached by two orthogonal massless springs to the origin of the coordinate system at the center post. Due to restoring forces proportional to the distance between the ball and the center post, the ball was accelerated towards the center post ($x = 0; y = 0$)³. At time t , the equations for the position of the ball in x - and y -directions were:

$$\begin{aligned}x(t) &= A_x \sin(\omega t + \phi_x) e^{-\frac{t}{\tau}} \\y(t) &= A_y \sin(\omega t + \phi_y) e^{-\frac{t}{\tau}}\end{aligned}$$

The amplitudes A_x and A_y and the phases ϕ_x and ϕ_y of the sinusoidal motions of the two springs were calculated from the ball's x - y position and velocity at release, which were converted into angle and velocity with respect to the center post. The motions were lightly damped to approximate realistic behavior, with the parameter τ describing the relaxation time for the decaying trajectory (for more detail see Müller & Sternad, 2004).

The execution space with the solution manifold was calculated numerically in a forward manner using the spring model of the task. For each point in execution space, defined by angle and velocity at release, the ball trajectory was computed for 2/3 of one traversal (the elliptic trajectory returned to its release position with slight deviation due to damping). For each point on the trajectory the distance to the target was evaluated. The point with the minimum distance to the target center was chosen as the result value. The result values were mapped onto gray shades and displayed as shown in Figure 1B. The solution manifold is defined as the subset of solutions with zero error. For better visualization, the Figure 1B shows those solutions in white that have results smaller than 1.2 cm.

Analysis of T -, N -, and C -Costs

To extract the three components for a set of data, 60 consecutive trials, each described by two execution variables (release angle and release velocity) and one result variable (distance to target or error), were analyzed as a set. Trials with a performance error more than three standard deviations different from the mean of that set were removed. There were, on average, 1.5 of these per set of 60 trials. For each set of 60 trials the mean result determined from all 60 trials was calculated. For the calculation of each component, the data were transformed in a specific way to create another set with optimal results in terms of one component while other features

of the data set remained unchanged. The mean result for this optimized data set was calculated. The algebraic difference between the mean result from the actual data set and the mean result from the optimized data set expressed the cost of the specific component. Figure 3 shows actual and optimized data sets for one subject from three different phases of practice. Each row shows results from a different day of practice. Each column shows data optimized in terms of a different component.

T-Cost expresses the cost to performance of a given data set not being centered in the best region of execution space. To estimate *T-Cost* for a data set, an optimized data set was generated in which the mean of the release angles and the mean of the release velocities were shifted to the location in execution space that yielded the best overall result. Importantly, the dispersion along each axis was preserved. This procedure is equivalent to optimizing the location without changing the noise and covariation of the data set. In the numerical procedure, the data was shifted on a grid of 600×600 possible center points for the data set. The boundary of this grid was determined by the limits of the task. The angles tested as centers were limited to those between 0 and -180 deg. This range was determined by the instructions given to the participants and the range of motion of the elbow joint. The velocities tested as centers were those between 200 and 800 deg/s. This range was determined by the solution manifold (there were no good solutions at lower release velocities) and by the velocities attainable in this task. The optimization procedure moved the whole data set through the grid and evaluated its mean result at each location. For data points that extended beyond the grid, the values were calculated on the extended execution space. The location that yielded the best overall performance result was compared with the result of the actual data set. The algebraic difference between the actual result and the result at the optimized center defined *T-Cost*. The panels in the leftmost column show the actual and optimized data sets used to calculate *T-Cost* for one subject at three phases of learning. The numbers in the top of each panel show the exact value for the example data sets. The three panels show how at three stages of practice *T-Cost* becomes increasingly smaller.

N-Cost—To estimate the cost to performance of a given data set that does not have optimal noise in execution variables, a virtual data set was created where variability was reduced to achieve the best possible mean result; the pairings of the angles and velocities were unchanged and the mean of angle and velocity was also unchanged. A first expectation would be that the mean of the data set, a single point, has minimal noise and is therefore optimal. However, note that each data set is evaluated in terms of its result, the minimum distance to the target. Hence, the mean of a data set may not be at the point with the optimal result. We therefore adopted the following procedure to assess optimal noise: First the mean of the angle and velocity was determined. Then, the data set was “shrunk” in 100 steps to collapse onto the mean. For the shrinking procedure, first the radial distance of each point to the mean was determined. Second, this radial distance was divided into 100 steps, each bringing every data point 1% closer to the mean angle and mean velocity until all 60 points collapsed at the mean. At each step the overall mean result was evaluated. The step that led to the best overall performance was selected as the set with optimal noise. Note that this method allows for the possibility that the mean may not be located on the best possible point in execution space with the best result. Hence, some noise may be of benefit to the result. The algebraic difference between the mean result of the optimal data set and the mean result of the actual data set defined *N-Cost*. The three panels in the middle column show the actual and optimized data sets created to calculate *N-Cost* for one subject at three phases of learning. The numbers in the top of each panel show the exact values for each of the example data sets and illustrate how *N-Cost* decreases with advanced skill level.

C-Cost expresses the cost to performance of a given data set not fully exploiting redundancy in the execution space. To estimate *C-Cost* for a data set, an optimized data set was generated in which the means and distributions of the angles and velocities were maintained, while the

individual pairs were recombined to achieve the best possible performance. This idealized data set was found with a greedy hill climbing algorithm using a pairwise matching procedure (Russell & Norvig, 2002). To implement this, first the pairs of angles, a_i , and velocities, v_i , were rank ordered from best to worst according to the error result, r_i , with $i = 1, 2, 3 \dots 60$. Next, the angle from the worst performing pair a_{60} was paired with v_{59} , and a_{59} was paired with v_{60} , i.e., v_{60} and v_{59} were swapped; the mean result of r_{59} and r_{60} was determined. If the mean result improved over the original r_{59} and r_{60} , the swap was accepted. As a next step v_{60} was swapped with v_{58} and the resulting mean error of r_{58} and r_{60} was evaluated. If the mean result improved, the swap was accepted. This procedure continued until a_{60} was paired with v_1 , i.e., v_{60} was swapped with v_1 . After this sequence of 59 comparisons, the same procedure was repeated with a_{59} . Hence, the batch consisted of $59 \times 59 = 3481$ comparisons. The number of profitable swaps was recorded for each batch. Then, this entire batch of procedures was repeated on the improved set until the number of swaps converged to zero (no further swaps could profitably be made).

The algebraic difference between the mean result of the actual data set and the mean result of the optimally recombined set defined *C-Cost*. Note that this method can find optimal covariation between execution variables even in situations in which the solution manifold is nonlinear, as in this task. Hence it is different from the statistical covariance calculation, which is essentially linear and applied to two variables. Importantly, the present calculation can also be applied to more than two variables, although it becomes numerically more involved (see discussion below). The three panels in the right column show the actual and optimized data sets used to calculate *C-Cost* for one subject at three phases of learning. Again, the numbers in the top of each panel show the exact values for each of the example data sets. They decrease with progressing experience as covariation becomes more pronounced.

Results

Result Variable: Error

Decrease of error and variability with practice is a robust phenomenon, frequently documented in the motor learning literature. In order to attest that our data indeed show these signatures of performance improvement and are consistent with many other motor learning studies, a first analysis evaluated the traditional error measures. Figure 4 displays the means and standard deviations of error over blocks of trials for the group of 12 subjects and, separately, for the expert subjects who practiced for 15 days. The error bars denote the standard error of the mean across the group of participants. As expected, there was a decrease across practice in both measures, with the greatest changes at the beginning. The experts' errors and standard deviations were lower than those of the average participants from beginning to end. The expert group continued to show improvement in both measures through block 45. Exponential functions highlight these changes over time. Note that the first value was excluded from the fitting procedure to avoid distortion of the fits in later stages of practice due to this generally very high initial value. The fitting parameters and the goodness of the fits are reported in Table 1.

Execution Variables: Release Angle and Velocity

Figure 5 displays the changes in release angle and release velocity over blocks of trials for the average subjects and the expert subjects. The error bars denote the standard error of the mean across participants. Panel A shows mean release angle over blocks, and Panel B shows mean release velocity over blocks. The angle variable tended to stabilize after block 9 (third day of practice). The velocity variable showed some gradual change towards higher values in the expert subjects. Panel C shows the standard deviations of release angle per block, and Panel D shows the standard deviation of release velocity per block. As these variability measures

decrease with practice in a similar fashion to the error measures, exponential functions highlight the decline (see Table 1). While the standard deviations of both variables indicated a plateau in the average participants, they continued to decrease after 6 days in the expert subjects.

Quantification of Costs

Figures 6 and 7 depict the amount of error accounted for by each of the three components, i.e., the cost to performance. In Figure 6 the means and the standard errors for both participant groups are illustrated. The decline across blocks of practice is highlighted by fitted exponential functions (excluding the first data point); the fitting parameters are included in Table 1. *T-Cost* tends to decrease significantly from block 1 to block 2 and reaches a low level close to zero very early in practice. *N-Cost* is comparable for both groups but shows pronounced decrease in the expert group throughout the 15 days. *C-Cost* in the average group plateaus while the experts showed continued decrease throughout practice. Figure 7 shows the exponential fits of all three components for ease of comparison. Panel A shows the average costs for the 12 average participants, Panel B, C, and D shows the three costs for each of the three experts. Figure 8 summarizes the results of the 12 individual participants in the average group by showing the rank ordering of the three costs in each block. Panel A presents the number of participants for whom each component made the largest contribution to error across blocks of practice. The figure illustrates the number of occurrences across the 12 participants for which each of the three costs had rank 1 in the respective block. Panel B shows the number of participants for whom each component made the smallest contribution to error across blocks of practice. The figure illustrates the number of occurrences across the 12 participants for which each of the three costs had rank 3 in the respective block.

T-Cost—Figure 6A shows the cost to performance of non-optimal tolerance. For most participants, *T-Cost* was the first component to be maximally exploited (Figure 7A), dropping rapidly from relatively high values in the first blocks to very low values. It ranked as the largest contributor to cost for 5 subjects in the first block, and it made the smallest contribution to error for 11 of the subjects by the sixth block (Figure 8). The expert participants showed the same trend, with *T-Cost* reaching values close to zero after block 20 (Figure 6A).

N-Cost—Figure 6B shows the cost to performance of non-optimal dispersion of the execution variables. *N-Cost* was the largest contributor to error overall as evident from Figure 7. It was the highest contributor for 7 participants in the first block and for 10 participants in the last block (Figure 8). The expert participants began with a lower *N-Cost* than the group average and continued improving for longer to reach similar values as *C-Cost*.

C-Cost—Figure 6C shows the cost to performance of non-optimal covariation. For the average participants, *C-Cost* declined relatively slowly across all days and blocks of practice, and the decline for the experts was similar. *C-Cost* was generally higher than *T-Cost* but lower than *N-Cost* (Figure 7). This general impression is mirrored in the rank analysis of the individual subjects (Figure 8).

Discussion

There were two major goals to this study. The first was to take the concepts of tolerance, noise, and covariation, first proposed by Müller & Sternad (2003, Müller & Sternad (2004) to quantify components of skill improvement, and to formulate them in terms of costs with respect to optimal performance. The second goal was to apply the new cost-analysis to investigate the relative importance of different costs and their changes during six and as well 15 days of practice in which the skill was learned and fine-tuned.

TNC-Cost Analysis and Comparison to Related Methods

Both the ΔTNC -analysis developed by Müller and Sternad and the *TNC-Cost* analysis developed here (together referred to as the *TNC*-approach) are related to but also distinct from other methods in motor control that evaluate distributional properties in data sets (Cusumano & Cesari, 2006; Scholz & Schöner, 1999). A central feature of the *TNC*-approach is the known mapping between a complete set of execution variables and their corresponding results as specified by the task goal. On the basis of this functional relation, analyses can be conducted on distributional properties of execution variables and their effects on results. In contrast to extant methods such as the UCM- and Cusumano and Cesari's approach that capitalize on the well-established mathematical methods of null space analysis, the *TNC-Cost* analysis developed new numerical methods to quantify three components. The UCM-approach analyzes the geometric properties of a data set by projecting the data onto the tangent space around a point (typically the mean of the data). This mathematical approach for detecting linear covariance among variables, which is mathematically related to principal component analysis, has the advantage of relying on established tools, but it also has some practical disadvantages. For instance, the linearization is typically performed at the mean of the data set, which may have a constant bias from the desired result. In contrast, the *TNC*-approach evaluates relations among execution variables in terms of their effects on the result. Further, *TNC-Cost* calculations do not rely on linearization, but can quantify alignment of data along any nonlinear manifold, such as the highly nonlinear one employed in this study (Müller & Sternad, 2003, 2008).

Another core feature of the *TNC*-approach that differs significantly from other methods is the assessment of tolerance or sensitivity. Standard sensitivity analysis is based on the Jacobian and evaluates the sensitivity of a point with respect to infinitesimally small perturbations. While this has the advantage of applying proven mathematical tools, assessing the response to infinitesimally small perturbations does not take into account that there may be a sizable dispersion around a point. Given the different layout of execution space and different curvature at different locations, a single solution in execution space may be locally stable, but an extended neighborhood can make a solution more or less robust with respect to variability. Further, a discontinuity in a more extended neighborhood can have severe negative effects on the result that go unnoticed in a local sensitivity analysis. In response to these limitations of standard sensitivity analysis, tolerance or *T-Cost* assesses the effect of dispersion in the execution variables on result over an extended neighborhood. Importantly, this neighborhood is defined by the dispersion and covariation of the data set. Hence, *T-Cost* is co-defined by the *N-Cost* and *C-Cost*. For example, a large dispersion may favor a location with a low curvature, while a reduced dispersion may favor a different location. This highlights the important point that variability may determine the strategy that subjects choose.

C-Cost as a measure of covariation is also significantly different from covariance, which is central to principal component analysis, and is also applied by UCM, and the approach developed by Cusumano and Cesari (2006). Covariance is essentially a linear concept measured in units of real-valued random variables. In contrast, *C-Cost* is measured in units of the result, i.e., centimeters in the present case. Therefore, the values have physical meaning and are directly interpretable. Further, covariation and *C-Cost* can be evaluated for more than two variables, and these variables can be nonlinearly related. Note that we previously estimated covariation in comparison to a permuted data set that had no covariation (Müller & Sternad, 2003). That quantification is conceptually different from the one proposed here that compares actual performance to that with the best possible covariation.

One last point to highlight is that the *TNC*-approach does not require that the execution variables have the same units, because comparison between different solutions is performed in the space and units of the result. Hence, the execution space does not need to have a metric. Another

aspect inherent to the two-level approach of *TNC* is that the functional relation between result and execution predicts the outcome of all execution strategies, which allows a priori predictions about successful and preferred movement strategies. For example, recent work by our group tested the prediction that the most tolerant solutions would be preferred by actors, a prediction that was supported for two task variations with different solution manifolds (Hu et al., submitted).

The rationale and results of our study are also in overall accordance with a series of experiments by Trommershäuser and colleagues that tested predictions from a decision-theoretic formulation of the problem (Trommershäuser, Gepshtein, Maloney, Landy, & Banks, 2005; Trommershäuser, Maloney, & Landy, 2003). Using a speed-accuracy task where subjects point to a target area that is bounded by a penalty area (at different distances and with different penalties), the distribution of hits is examined with respect to the expected gain (reward and penalty). Importantly, and different from our approach, hitting success was binary (positive for the target area and negative for the penalty area), and subjects maximized the cumulative score over series of trials. Formalized in a decision theoretical framework where a gain function is optimized based on the weighted sum of the gain and the subject's inherent variability, the results showed systematic effects of the location and relative magnitude of the penalty and the target area on the distributions of end-point. The results supported the idea that selection of a movement strategy is determined by the subject's inherent variability. In contrast to our approach, their work places emphasis on selection of strategies, not learning, and does not differentiate between routes participants might take to optimize their score.

TNC-Cost* Analysis and Comparison to the Previous *TNC-Decomposition

Although the concepts of tolerance, noise, and covariation were drawn from Müller and Sternad's earlier work, there are two major differences in the way they are conceived and implemented here: (1) comparisons to idealized data sets; and (2) independent calculations of components. The first difference is more central and will be described first. In the previous *TNC*-method, the three components are calculated as differences between two blocks of data, for example between a block at the beginning of practice and a block later in practice. The changes in the three components, ΔT , ΔC , and ΔN , are reported as changes in performance. This has the disadvantage that the results do not give an absolute estimate of the contribution of a component but only a relative one, which makes interpretation of inter-individual differences difficult. In the *TNC-Cost* approach developed here, the reference set for each block of data is an optimized version of itself. Note that this requires a logical reconceptualization of the components from *benefits* to *costs*. Costs to performance are defined as the algebraic differences between the results of the actual data sets and those of the virtual data sets. Each cost is an absolute measure of how much this data set could improve if it optimized in terms of one component. The payoff from this change from a relative to an absolute estimate is that now the contributions can be quantified for a single set of data. Furthermore, meaningful comparisons can be made across subjects. For instance, in the present data the expert participants had lower *T-Costs* in the first block than the average participants. Similarly, the method could be used to make comparisons between healthy participants and participants with motor disorders.

Unlike in Müller and Sternad's (2004) original ΔTNC -analysis, the components in the *TNC-Cost* analysis are computed independently. Thus, the order of calculations does not matter, and no assumptions are required about priority or precedence of components. However, the price for relinquishing these assumptions is that components may overlap, which is not possible in the original procedure. That method is applied in a nested fashion, such that tolerance ΔT is estimated by the comparison of two data sets in which covariation has already been removed; thus, potential contributions from covariation are eliminated. Noise reduction ΔN is then

defined as the residual after ΔT and ΔC are subtracted from the total change in error. This has the advantage that the sum of the three components is necessarily equal to the error. In contrast, calculating the components independently allows overlap to occur in the costs and they do not sum to the actual error. As further detailed below, some caution is needed in their interpretation. To recapture the benefit of components that sum, the present method of comparison to idealized data sets could be similarly implemented in a nested way by first shifting the data to an ideal location, then arranging it to optimize covariation, and then reducing variability as needed to achieve perfect performance. To illustrate the differences between the two methods, we conducted both analyses on one participant's data (Expert 3). The results of the ΔTNC method are plotted in a cumulative fashion across blocks to convey the gradual improvement. Due to the conceptual differences the results should not be compared directly.

The inter-relation between components is an important aspect of the new method and some caveats have to be made clear. As stated before, in the creation of each virtual data set, only one aspect of variability in execution variables is optimized. However, this optimization can affect more than one component. Importantly, the shape and orientation of the solution manifold at a given location determines the degree to which covariation can play a role. For example, in the limit case that the manifold is linear and parallel to one dimension, C -Cost is zero. Further, shifting a data set to a more tolerant location on the solution manifold may increase or decrease the amount of covariation and noise possible, and thus affect the magnitude of C -Cost and N -Cost.

Of particular importance for understanding the role of motor noise in performance is the fact that the magnitude of N -Cost depends both on T -Cost and on C -Cost. A given variance in the execution variables may be harmful in an error-sensitive region of the solution manifold but may have negligible effect on the result in a less error-sensitive, i.e., more tolerant region of the solution manifold. Similarly, a given dispersion in execution variables with poor covariation (not aligned with the solution manifold) may be harmful to performance, while the same variance at the same location but with good covariation (aligned with the solution manifold) may allow perfect performance. Understanding the relations among the components can provide interesting insights into changes in variability across practice. For instance, performers who have reached their physiological limit in terms of reducing motor noise can reduce the impact of that motor noise on their performance (N -Cost) by continuing to optimize T -Cost and C -Cost. Note that while most performers in the study reported here did not continue to refine tolerance after the first few days of practice, more improvement in that component was possible, as illustrated by the performance of the expert participants.

As a final comment on the method, it should be mentioned that the quantification of the components is, to some degree, dependent on the chosen coordinate system, i.e., the space of execution variables. This issue was discussed in two earlier articles (Müller, Frank, & Sternad, 2007; Smeets & Louw, 2007) and extended by Müller & Sternad (2008). As demonstrated, the ΔTNC method is dependent on the choice of coordinates. However, similar dependencies exist for other methods of analyzing variance, including the related approaches of UCM and optimal feedback control. In fact, this lack of invariance across coordinate transformations is also present in simple statistical measures such as means.

Application of the TNC -Cost Analysis to Learning and Fine-Tuning a Skill

In the virtual throwing task used here, all participants became more accurate and consistent in their throws over the course of the experiment. Six days of practice were more than sufficient for most participants to reach a plateau in performance. All three components contributed to error in performance, and all three costs decreased across practice. Generally, the rank order of the contributions of the three costs was consistent throughout the experimental sessions, with T -Cost high at the beginning and then dropping rapidly, contributing the least to error

after the first day or two. *N-Cost* contributed the most to error, and *C-Cost* was intermediate to the other two.

In the first day of practice, the main determinant of performance improvement was the discovery of a solution, as quantified by pronounced changes in *T-Cost*. Later in practice, reductions in *T*-, *C*- and *N-Costs* made comparable contributions to improvement. *T-Cost* can be thought of as the cost of not using the best strategy, i.e., not being in the right region of execution space. Thus, finding a drop earlier than in the other components is broadly consistent with the idea that skill acquisition occurs in stages, with the first stage being more cognitive or strategic where actors discover the lay-out of solutions and find “the ball bark” of the best solution (Fitts & Posner, 1967). The larger changes of tolerance at the beginning of practice are also consistent with the findings of Müller and Sternad (2004) despite the fact that the *TNC*-calculations were different. Following this drop in *T-Cost*, the other two costs gradually decreased but maintained their relative ordering. Notably, in the average group *N-Cost* remains highest throughout all days of practice. Inspection of individual participants’ data reveals that this general ordering is observed in many individuals.

Three participants were of particular interest as they had considerable expertise in real-world throwing and developed a personal ambition to achieve solid sequences of perfect hits in this experimental task. Right from the start their performance was better than the average group and they continued to improve long after the group’s average performance had leveled off. By the end of practice their error scores were below the threshold that signaled success 75% of the time. Additionally, all three costs descended to lower asymptotic values for these three participants than for the average group. In fact, *T-Cost* diminished to almost zero towards the end of practice, and *N-Cost* descended as much in days 8–15 as it did in days 2–8. In relative terms it is worth noting that late in practice *N-Cost* decreased to be of equal magnitude as *C-Cost*. Potentially, this could be viewed as a sign of fine-tuning a skill.

Motor noise is often shown to be signal-dependent (Harris & Wolpert, 1998). Thus, one plausible expectation is that participants would gravitate toward the regions of the solution space that allow them to throw with lower velocity, in order to reduce their motor noise. Notably, very few of the participants chose this strategy. In fact, the participant with the best performance gravitated toward a solution region with fairly high velocity. This finding is also consistent with results in the study by (Hu et al., submitted) that specifically tested this hypothesis but found no evidence for it. Thus, other means of reducing the negative effects of motor noise on performance had to be employed.

Taking the results of the two groups together what can be learned from these results? Clearly, variability is more than just noise. Reducing variability in performance can be achieved in several ways that go beyond simply reducing noise. One important aspect is to find a strategy or location in the space of execution variables that is insensitive to variations in execution. A second aspect is to covary the variables in ways that optimize the result, in other words, channel them into directions that do not matter to the result. Lastly, reduce the dispersion in the execution variables, although the importance of this latter route highly depends on the former two. Hence, increasing consistency in performance is achieved by at least three routes. Possibly, the tuning of covariation and tolerance are the routes that minimize the effect of the inevitable motor noise.

Summary and Conclusions

The underlying determinants responsible for performance improvement and learning are not yet well understood. The *TNC-Cost* analysis offers a new approach for uncovering structure in data. In the throwing task examined here, the results of this analysis quantified the first exploratory stage of practice. The results further highlighted the development of covariation

between execution variables. Lastly, they showed that noise remains high and seems to be tuned last. Taken together, these results indicate that skill acquisition consists largely of participants' increasing sensitivity to relatively subtle aspects of the execution space.

The *TNC*-approach is suitable for exploring any task in which a performance goal can be defined as redundantly dependent on several execution variables. It has been applied to a ball-bouncing task (Boonstra, Wei, & Sternad, submitted), to a darts task (Müller & Sternad, 2008), and, as in this paper, to the skittles task. The method would also be suitable for investigating a multi-joint pointing movement, a postural task, or any other task in which execution variables can be projected into result space. Only the functional relationship between execution variables and results must be known.

Finally, the *TNC-Cost* analysis allows comparisons across participants. This makes it suitable for application to patients with different kinds of motor disorders, as well as for development of training regimes for athletes. Future investigations should continue to explore interventions targeting each of the three components as potential routes to improvement.

Acknowledgments


This research was supported by grants from the National Science Foundation, BCS-0450218, the National Institutes of Health, R01 HD045639, and the Office of Naval Research, N00014-05-1-0844. We would like to express our thanks to Hermann Müller, Tjeerd Dijkstra and Neville Hogan for many insightful discussions and to Kunlin Wei for his help in creating the experimental set-up.

References

- Beek PJ, van Santvoord AAM. Learning the cascade juggle: A dynamical systems analysis. *Journal of Motor Behavior* 1992;24(1):85–94. [PubMed: 14766500]
- Bernstein, N. The coordination and regulation of movement. London: Pergamon Press; 1967.
- Bernstein, N. Dexterity and its development. In: Latash, ML.; Turvey, MT., editors. Dexterity and its development. Champaign: Human Kinetics; 1996.
- Burton RM, Mpitsos GJ. Event-dependent control of noise enhances learning in neural networks. *Neural Networks* 1992;5:627–637.
- Churchland MM, Afshar A, Shenoy KV. A central source of movement variability. *Neuron* 2006;52:1085–1096. [PubMed: 17178410]
- Craig, JJ. Introduction to robotics. Reading, MA: Addison-Wesley; 1986.
- Cusumano JP, Cesari P. Body-goal variability mapping in an aiming task. *Biological Cybernetics* 2006;94(5):367–379. [PubMed: 16501988]
- Duarte M, Zatsiorsky VM. Patterns of center of pressure migration during prolonged unconstrained standing. *Motor Control* 1999;3(1):12–27. [PubMed: 9924098]
- Fitts PM. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology* 1954;47:381–391. [PubMed: 13174710]
- Fitts, PM.; Posner, MI. Human performance. Belmont: Brooks/Cole; 1967.
- Full RJ, Kubow T, Schmitt J, Holmes PJ, Koditschek DE. Quantifying dynamic stability and maneuverability in legged locomotion. *Integrative and Comparative Biology* 2002;42:149–157.
- Geman S, Bienenstock E, Doursat R. Neural networks and the bias/variance dilemma. *Neural Computation* 1992;4:1–58.
- Harris CM, Wolpert DM. Signal-dependent noise determines motor planning. *Nature* 1998;394:780–784. [PubMed: 9723616]
- Hasan Z. The human motor control system's response to mechanical perturbation: Should it, can it, and does it ensure stability? *Journal of Motor Behavior* 2005;37(6):484–493. [PubMed: 16280319]
- Hu X, Müller H, Sternad D. Neuromotor noise and sensitivity to error in the control of a throwing task. submitted
- James, RC.; James, G. Mathematics dictionary. New York: Chapman and Hall; 1992.

- Jones KE, Hamilton AF, Wolpert DM. Sources of signal-dependent noise during isometric force production. *Journal of Neurophysiology* 2001;88:1533–1544. [PubMed: 12205173]
- Kudo K, Tsutsui S, Ishikura T, Ito T, Yamamoto Y. Compensatory correlation of release parameters in a throwing task. *Journal of Motor Behavior* 2000;32(4):337–345. [PubMed: 11114227]
- Lashley KS. Basic neural mechanisms in behavior. *The Psychological Review* 1930;37(1):1–24.
- Latash ML, Scholz JF, Danion F, Schöner G. Structure of motor variability in marginally redundant multifinger force production tasks. *Experimental Brain Research* 2001;141(2):153–165.
- Liégeois A. Automatic supervisory control of the configuration and behavior of multibody mechanisms. *IEEE Transactions on Systems, Man, and Cybernetics, SMC-7* 1977;(12):868–871.
- Martin TA, Gregor BE, Norris NA, Thach WT. Throwing accuracy in the vertical direction during prism adaptation: not simply timing of ball release. *Journal of Neurophysiology* 2001;85:2298–2302. [PubMed: 11353043]
- Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E. Equations of state calculations by fast computing machines. *Journal of Chemical Physics* 1953;21(6):1087–1092.
- Müller H, Frank T, Sternad D. Variability, covariation and invariance with respect to coordinate systems in motor control. *Journal of Experimental Psychology: Human Perception and Performance* 2007;33(1):250–255. [PubMed: 17311492]
- Müller H, Sternad D. A randomization method for the calculation of covariation in multiple nonlinear relations: Illustrated at the example of goal-directed movements. *Biological Cybernetics* 2003;39:22–33.
- Müller H, Sternad D. Accuracy and variability in goal-oriented movements: decomposition of gender differences in children. *Journal of Human Kinetics* 2004a;12:31–50.
- Müller H, Sternad D. Decomposition of variability in the execution of goal-oriented tasks – Three components of skill improvement. *Journal of Experimental Psychology: Human Perception and Performance* 2004b;30(1):212–233.
- Müller, H.; Sternad, D. Motor learning: Changes in the structure of variability in a redundant task. In: Sternad, D., editor. *Progress in Motor Control - A Multidisciplinary Perspective*. New York: Springer; 2008.
- Newell, KM.; McDonald, PV. Searching for solutions to the coordination function: Learning as exploratory behavior. In: Stelmach, GE.; Requin, J., editors. *Tutorials in motor behavior II*. Amsterdam: Elsevier; 1992. p. 517-532.
- Ölveczky BP, Andalman AS, Fee MS. Vocal experimentation in the juvenile songbird requires a basal ganglia circuit. *PLOS Biology* 2005;3(5):e153. [PubMed: 15826219]
- Riccio, GE. Information in movement variability about the quantitative dynamics of posture and orientation. In: Newell, KM.; Corcos, DM., editors. *Variability and motor control*. Champaign, IL: Human Kinetics; 1993. p. 317-357.
- Russell, S.; Norvig, P. *Artificial intelligence*. New York: Prentice Hall; 2002.
- Schmidt RA, Zelaznik HN, Hawkins B, Frank JS, Quinn JT. Motor-output variability: a theory for the accuracy of rapid motor acts. *Psychological Review* 1979;47(5):415–451. [PubMed: 504536]
- Scholz JP, Schöner G. The uncontrolled manifold concept: identifying control variables for a functional task. *Experimental Brain Research* 1999;126:289–306.
- Scholz JP, Schöner G, Latash ML. Identifying the control structure of multijoint coordination during pistol shooting. *Experimental Brain Research* 2000;135:382–404.
- Smeets JBJ, Louw S. The contribution of covariation to skill improvement is an ambiguous measure. *Journal of Experimental Psychology: Human Perception and Performance* 2007;33(1):246–249. [PubMed: 17311491]
- Stimpel, E. Der Wurf (On throwing). In: Krüger, F.; Klemm, O., editors. *Motorik (On motor control)*. München, Germany: Beck; 1933. p. 109-138.
- Trommershäuser J, Gepshtein S, Maloney LT, Landy MS, Banks MS. Optimal compensation for changes in task-relevant movement variability. *Journal of Neuroscience* 2005;25(31):7169–7178. [PubMed: 16079399]
- Trommershäuser J, Maloney LT, Landy MS. Statistical decision theory and trade-offs in the control of motor responses. *Spatial Vision* 2003;16(3–4):255–275. [PubMed: 12858951]

van Beers RJ, Haggard P, Wolpert DM. The role of execution noise in movement variability. *Journal of Neurophysiology* 2004;91:1050–1063. [PubMed: 14561687]






Figure 1.

Work space and execution space for three hypothetical throws. **A:** Three exemplary throws in work space as participants see in the experiment. The view is a top-down view onto the pendular skittles task. **B:** The release variables of the same three throws represented in execution space. The grey shades indicate the level of success for different combinations of release angle and velocity. White denotes the set of release angles and velocities that lead to a direct hit, the solution manifold. A perfect hit is defined as the center of the ball passing within 1.3 cm of the center of the target.




Figure 2.
Experimental setup.

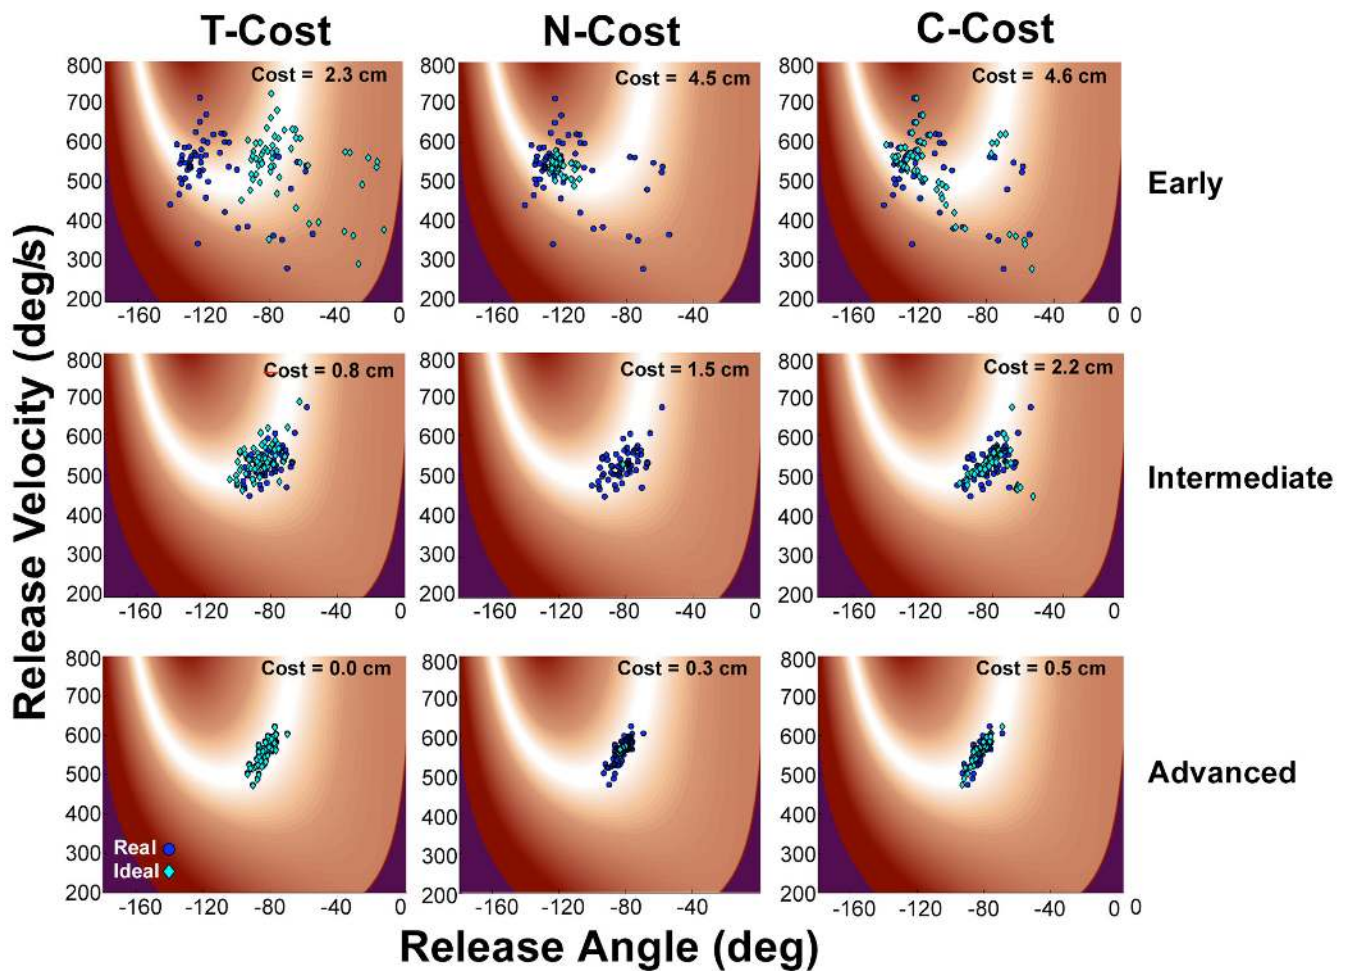


Figure 3.

Exemplary data sets and the corresponding virtual data sets used for the analysis of *T-Cost*, *C-Cost*, and *N-Cost*. Exemplary data from one participant are shown from three blocks. Circles represent actual throws made by the participant, and diamonds represent surrogate data with one component idealized. The panels in the top row show data from the first block of practice on the Day 1; the middle row shows data from the first block of practice on the Day 6, and the bottom row shows data from the first block of practice on Day 15. The left column shows data optimized in terms of tolerance; the middle column shows data optimized in terms of covariation, and the right column shows data optimized in terms of noise. For more detail see text.

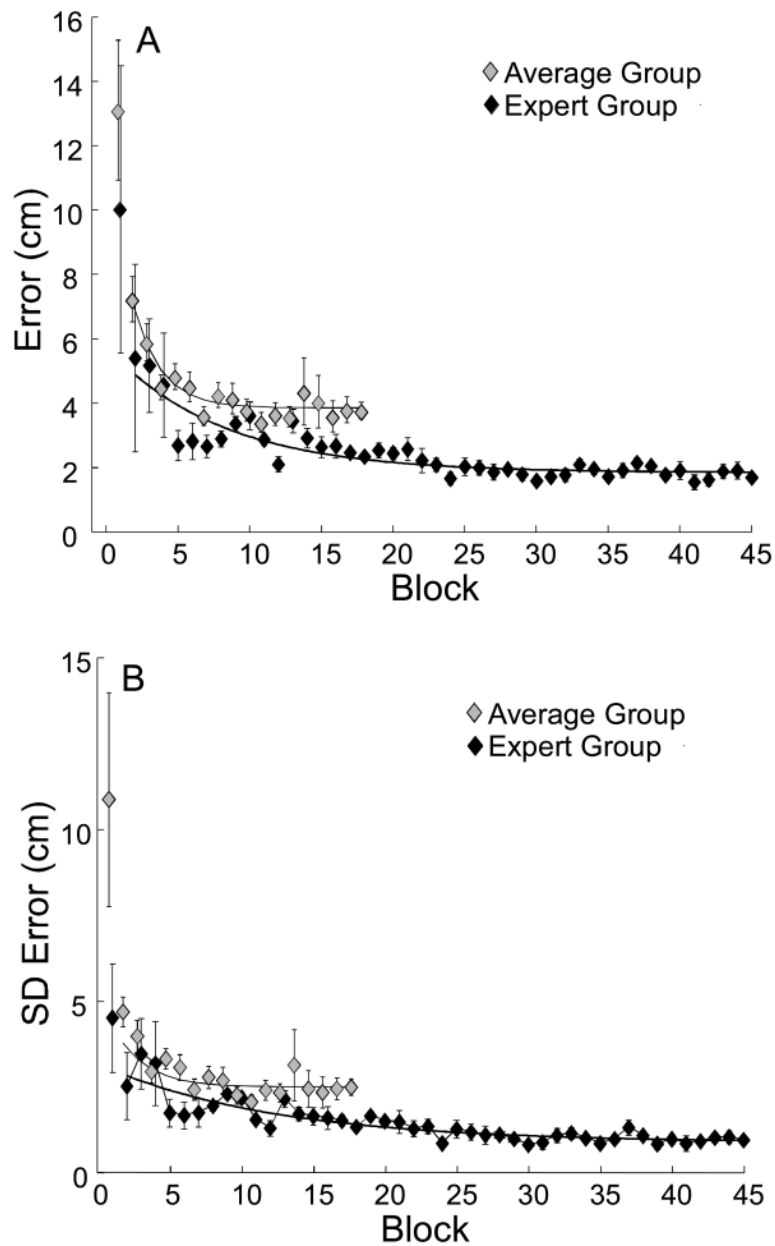


Figure 4. Changes in the error to the target with practice. For each block the data of 60 trials were averaged. Participants performed three blocks of trials per day. **A:** Mean error for each block of throws. Mean results for twelve participants over eighteen blocks (six days) of practice are shown in gray, and results for the expert participants over 45 blocks (15 days) of practice are shown in black. Error bars show standard error across participants. **B:** Standard deviations of error averaged over participants across blocks of trials. Error bars show standard error across participants.

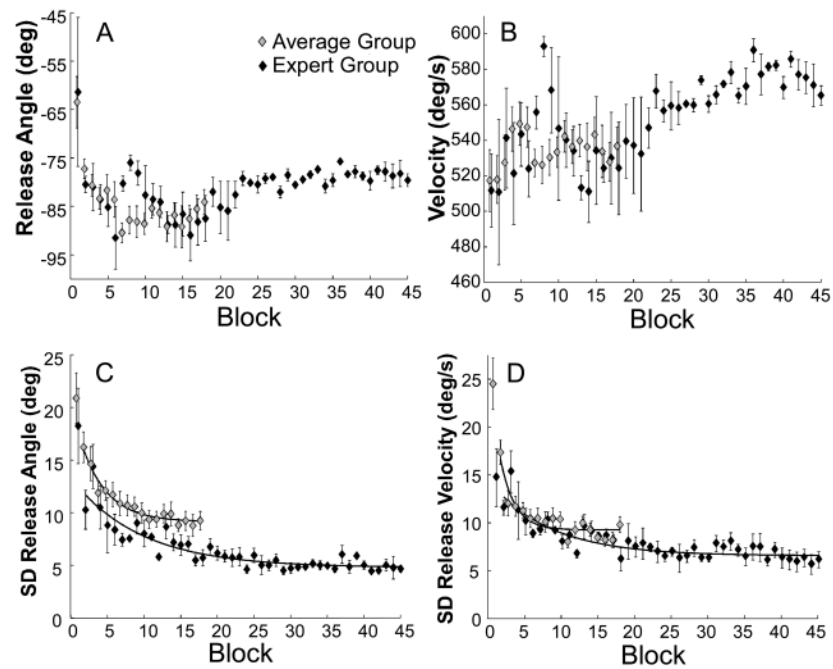


Figure 5. Changes of the two execution variables, angle and velocity at release, across blocks of practice. The gray symbols represent averages across 12 participants with the error bars denoting the standard error across participants. The black symbols denote the expert participants who practiced for 45 blocks (15 days). **A:** Average release angle. **B:** Average release velocity. **C:** Standard deviations of release angle. **D:** Standard deviations of release velocity.

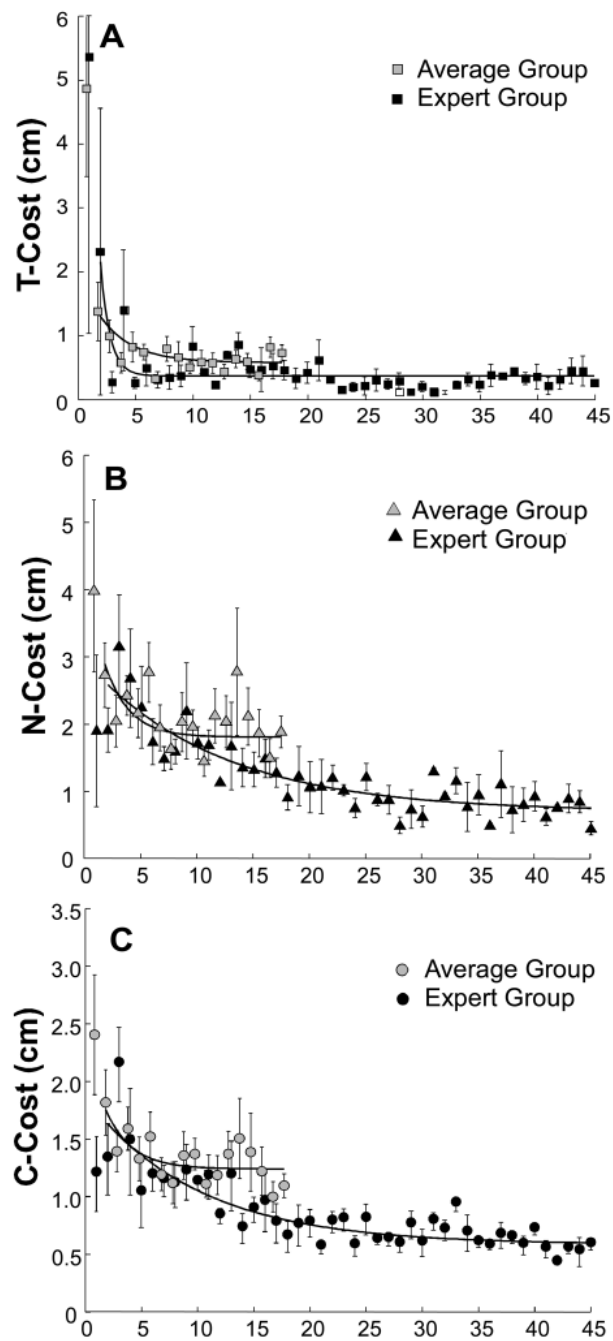


Figure 6. Contributions of *T-Cost*, *N-Cost* and *C-Cost* to error. Mean results for twelve participants over 18 blocks (six days) of practice are shown as in gray, and results for the expert group over 45 blocks (15 days) of practice are shown in black. Error bars show standard error across participants.

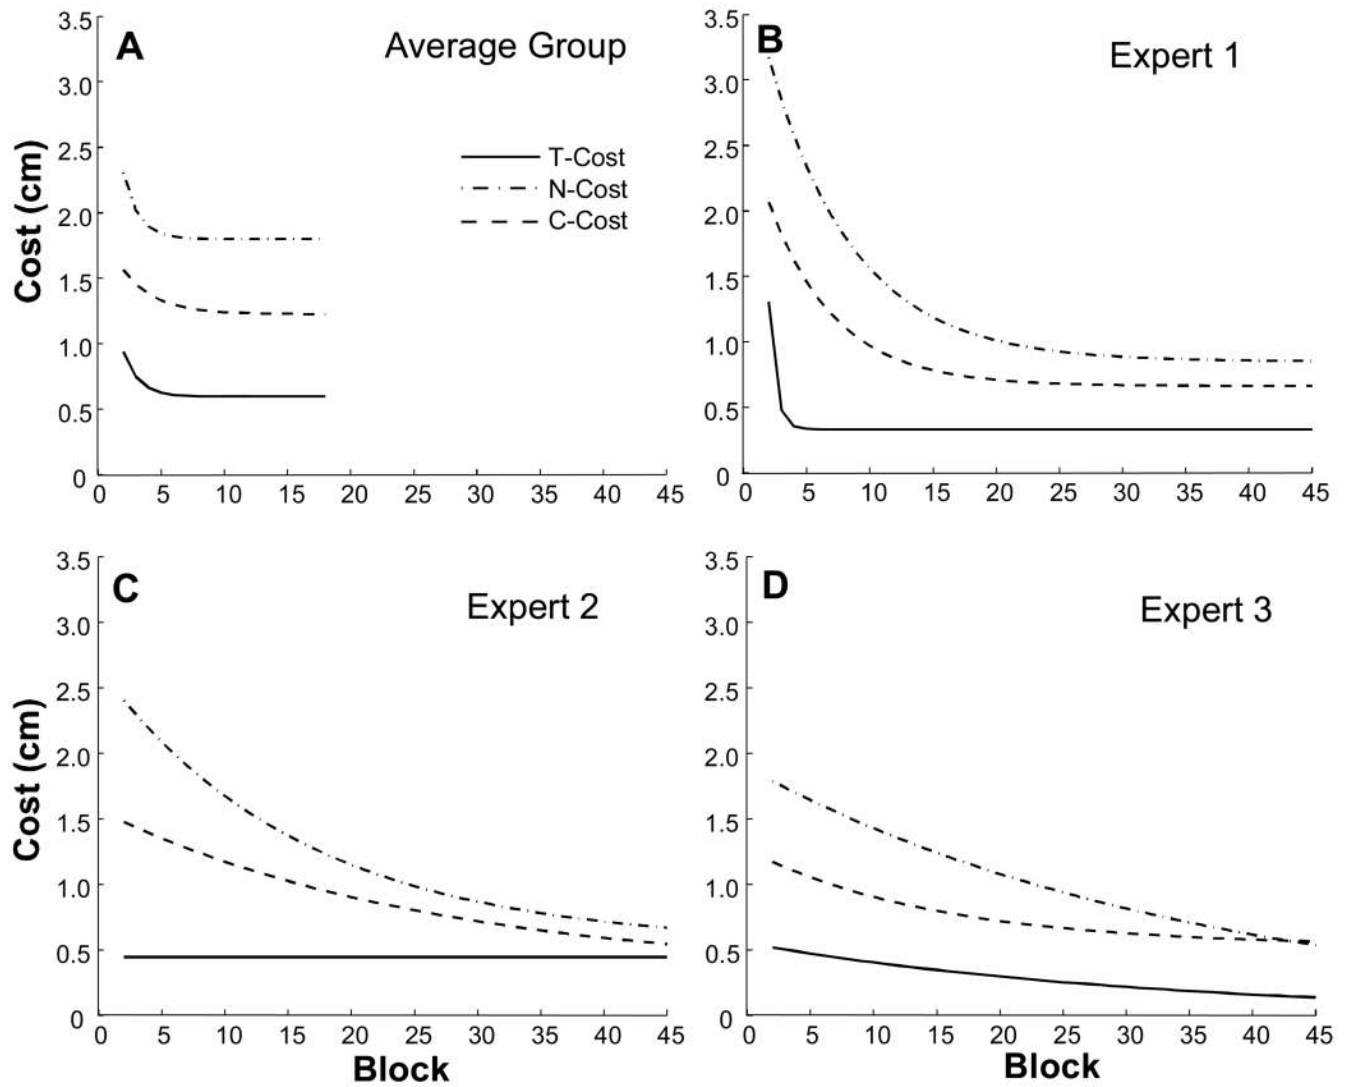


Figure 7. Exponential fits to *T-Cost*, *N-Cost*, and *C-Cost* for the average and the expert group data. **A:** Exponential fits to the three costs for the group. **B, C, D:** Exponential fits for the three expert participants.

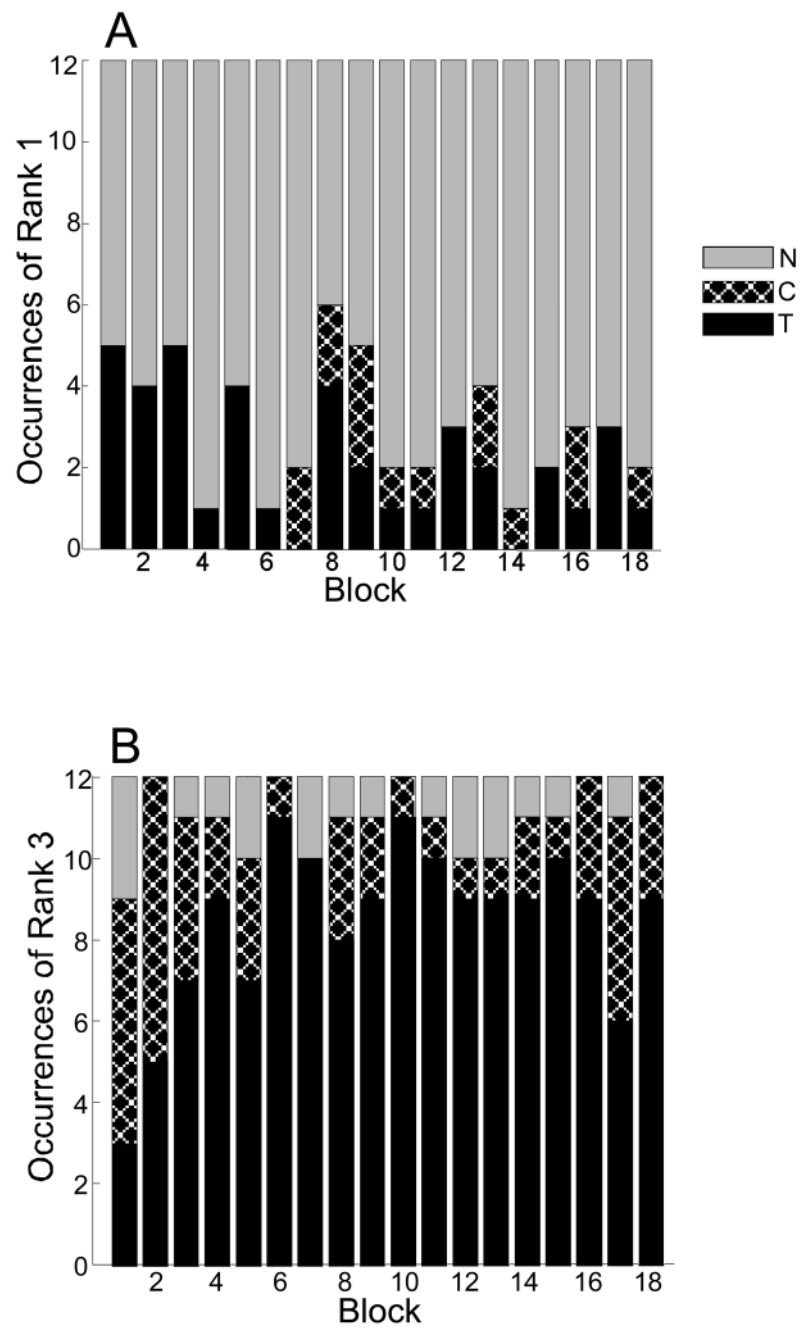


Figure 8. Summary of rank orders of the three costs across blocks and participants. **A:** Number of participants across practice blocks for whom each cost made the greatest contribution to error. **B:** Number of participants across blocks for whom each cost made the smallest contribution to error.

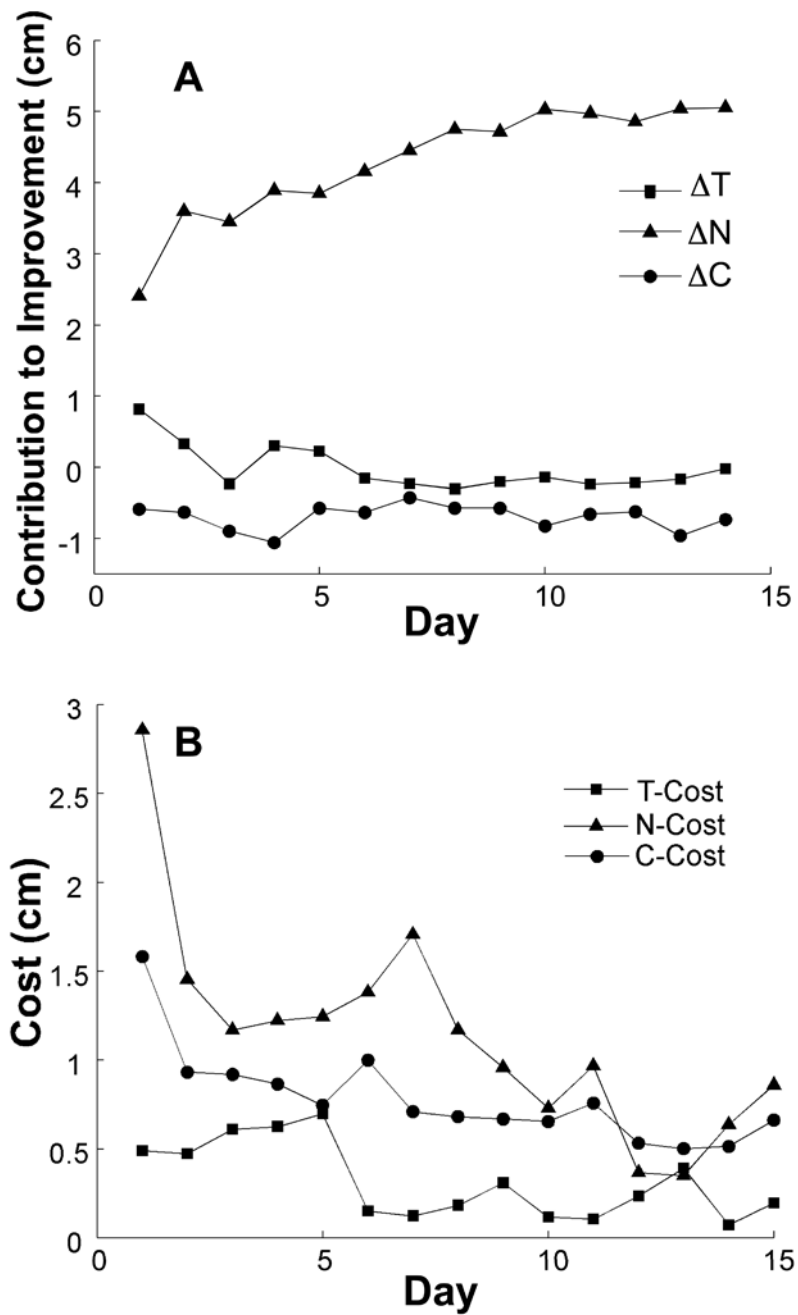


Figure 9. Results of the *TNC*-analyses applying the ΔTNC calculations (panel A) and the *TNC-Cost* calculations (panel B).

Table 1

Parameters and R^2 -values of the exponential fits to the group data of the average subjects and the expert subjects.

	Average Group			Expert Group			
	a	b	c	a	b	c	
Error	58.0	-.53	37.3	.90	.345	18.4	.79
SD Error	35.1	-.45	25.1	.84	21.1	8.8	.77
SD Release Angle	13.3	-.24	8.7	.96	7.7	4.8	.82
SD Release Velocity	209.5	-.52	94.1	.85	69.1	66.0	.75
T-Cost	15.8	-.35	5.8	.63	60.0	3.6	.59
C-Cost	12.1	-.43	12.3	.48	11.6	6.1	.76
N-Cost	27.5	-.46	18.0	.40	20.5	6.9	.78