

RESEARCH ARTICLE

A Hessian-based decomposition characterizes how performance in complex motor skills depends on individual strategy and variability

Paolo Tommasino^{1*}, Antonella Maselli^{1,2}, Domenico Campolo³,
Francesco Lacquaniti^{1,4}, Andrea d'Avella^{1,5*}

1 Laboratory of Neuromotor Physiology, IRCCS Fondazione Santa Lucia, Rome, Italy, **2** Institute of Cognitive Sciences and Technologies, National Research Council, Rome, Italy, **3** Synergy Lab, Robotics Research Centre, School of Mechanical and Aerospace Engineering, Nanyang Technological University, Singapore, **4** Department of Systems Medicine and Center of Space Biomedicine, University of Rome Tor Vergata, Rome, Italy, **5** Department of Biomedical and Dental Sciences and Morphofunctional Imaging, University of Messina, Messina, Italy

* tommasinopaolo@gmail.com (PT); a.davella@hsantalucia.it (AdA)



OPEN ACCESS

Citation: Tommasino P, Maselli A, Campolo D, Lacquaniti F, d'Avella A (2021) A Hessian-based decomposition characterizes how performance in complex motor skills depends on individual strategy and variability. PLoS ONE 16(6): e0253626. <https://doi.org/10.1371/journal.pone.0253626>

Editor: Welber Marinovic, Curtin University, AUSTRALIA

Received: June 1, 2021

Accepted: June 10, 2021

Published: June 30, 2021

Peer Review History: PLOS recognizes the benefits of transparency in the peer review process; therefore, we enable the publication of all of the content of peer review and author responses alongside final, published articles. The editorial history of this article is available here: <https://doi.org/10.1371/journal.pone.0253626>

Copyright: © 2021 Tommasino et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant data are within the paper and its [Supporting information files](#).

Abstract

In complex real-life motor skills such as unconstrained throwing, performance depends on how accurate is on average the outcome of noisy, high-dimensional, and redundant actions. What characteristics of the action distribution relate to performance and how different individuals select specific action distributions are key questions in motor control. Previous computational approaches have highlighted that variability along the directions of first order derivatives of the action-to-outcome mapping affects performance the most, that different mean actions may be associated to regions of the actions space with different sensitivity to noise, and that action covariation in addition to noise magnitude matters. However, a method to relate individual high-dimensional action distribution and performance is still missing. Here we introduce a decomposition of performance into a small set of indicators that compactly and directly characterize the key performance-related features of the distribution of high-dimensional redundant actions. Central to the method is the observation that, if performance is quantified as a mean score, the Hessian (second order derivatives) of the action-to-score function determines how the noise of the action distribution affects performance. We can then approximate the mean score as the sum of the score of the mean action and a tolerance-variability index which depends on both Hessian and action covariance. Such index can be expressed as the product of three terms capturing noise magnitude, noise sensitivity, and alignment of the most variable and most noise sensitive directions. We apply this method to the analysis of unconstrained throwing actions by non-expert participants and show that, consistently across four different throwing targets, each participant shows a specific selection of mean action score and tolerance-variability index as well as specific selection of noise magnitude and alignment indicators. Thus, participants with different strategies may display the same performance because they can trade off sub-optimal mean action for better tolerance-variability and higher action variability for better alignment with more tolerant directions in action space.

Funding: European Union's Horizon 2020 research and innovation programme under grant agreement No 644727 to AdA (<https://cogimon.eu/cognitive-interaction-motion-cogimon>); Italian Ministry of Health (IRCCS Fondazione Santa Lucia - Ricerca Corrente) to FL (http://www.salute.gov.it/portale/temi/p2_5.jsp?lingua=italiano&area=Ricerca%20sanitaria&menu=corrente) Italian Space Agency (grants I/006/06/0 and 2019-11-U.0) to FL (<https://www.asi.it/>) Italian University Ministry (PRIN grants 2015HFWRY to AdA and FL; 2017CBF8NJ_005 to FL; <https://prin.miur.it/index.php?pag=2015>) The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

1 Introduction

In many goal-directed human behaviors, such as throwing a projectile towards the center of a target, performance depends on how accurate is the outcome of repeated actions. In throwing tasks, *accuracy* of a single throw may be quantified by a *score*, e.g. a scalar function that penalizes/rewards motor outcomes depending on their distance from the desired target position [1, 2]. In this perspective, the goal of a thrower would be that of minimizing/maximizing the mean score over repeated trials [3, 4]. Because of bias and noise in the sensorimotor transformations mapping goals into actions [5, 6], the score typically varies across trials and hence the *performance* of a throwing strategy, defined as its mean score, will in general depend on the *distribution* of motor actions [1, 7].

In many tasks, the relationship between motor actions and their outcomes is *redundant* [8] and to different actions there might correspond the same task outcome, hence the same score. As an example, consider the throwing task shown in Fig 1A where the outcome space is the two-dimensional space of all the possible landing positions of the ball on a vertical board, and the score depends on where the ball lands with respect to the aimed target. The landing position of the ball ultimately only depends on the position and velocity with which the ball is released (*action variables*). The center of the target then can be hit with different combinations (or *covariations*) of such action variables: for instance one can hit the target by releasing the ball from different positions modulating the velocity vector accordingly, or from a fixed position but with different combinations of vertical and horizontal velocities. These different but *task-equivalent* actions resulting in the same landing position form a subset of the action space which is called *solution manifold*. Key questions in human motor control are then whether different individuals select specific action distributions to achieve a given performance level, what characteristics of the action distribution relate to performance, and how action distributions change when performance improves with practice [1, 8–12].

In well-developed motor skills, outcomes are typically unbiased (zero mean error) and hence outcome variability (or precision) is usually taken as a measure of performance. Indeed, when learning a new motor skill involving redundant actions with different amounts of noise or noise tolerance in different regions of the action space, participants improve performance by selecting actions whose outcome is less affected by motor noise [13, 14]. The relationship between action distribution and outcome variability in goal-directed behaviors has been addressed by computational approaches that take into account the geometry of the mapping between actions and outcomes, also known as the goal function, near the solution manifold. Methods such as the Uncontrolled Manifold (UCM) [15, 16] and the Goal-Equivalent Manifold (GEM) [17, 18] typically approximate the (non-linear) action-to-outcome function with a (locally) linear map, which relates (stochastic) perturbations of the mean action to the precision (variance or covariance) of task outcomes. More specifically, the gradient (or Jacobian) of such mapping is employed to quantify action variability along *task-irrelevant* directions (directions parallel to the solution manifold) and *task-relevant* directions (directions orthogonal to the solution manifold). The UCM applied to reaching, pointing and throwing tasks [15, 16, 19] has shown that *covariation* between redundant actions is an important mechanism used by skilled performers to push motor variability along the task-irrelevant directions, hence increasing the precision of their task outcomes. Differently from UCM, which only quantifies motor strategies in terms of the *alignment* between action variability and task-relevant/task-irrelevant directions, the GEM approach takes also into account the *sensitivity* of the solution manifold to local perturbations. Then, different choices of mean actions may result in different amounts of outcome variability because of the specific alignment and sensitivity to the different mean

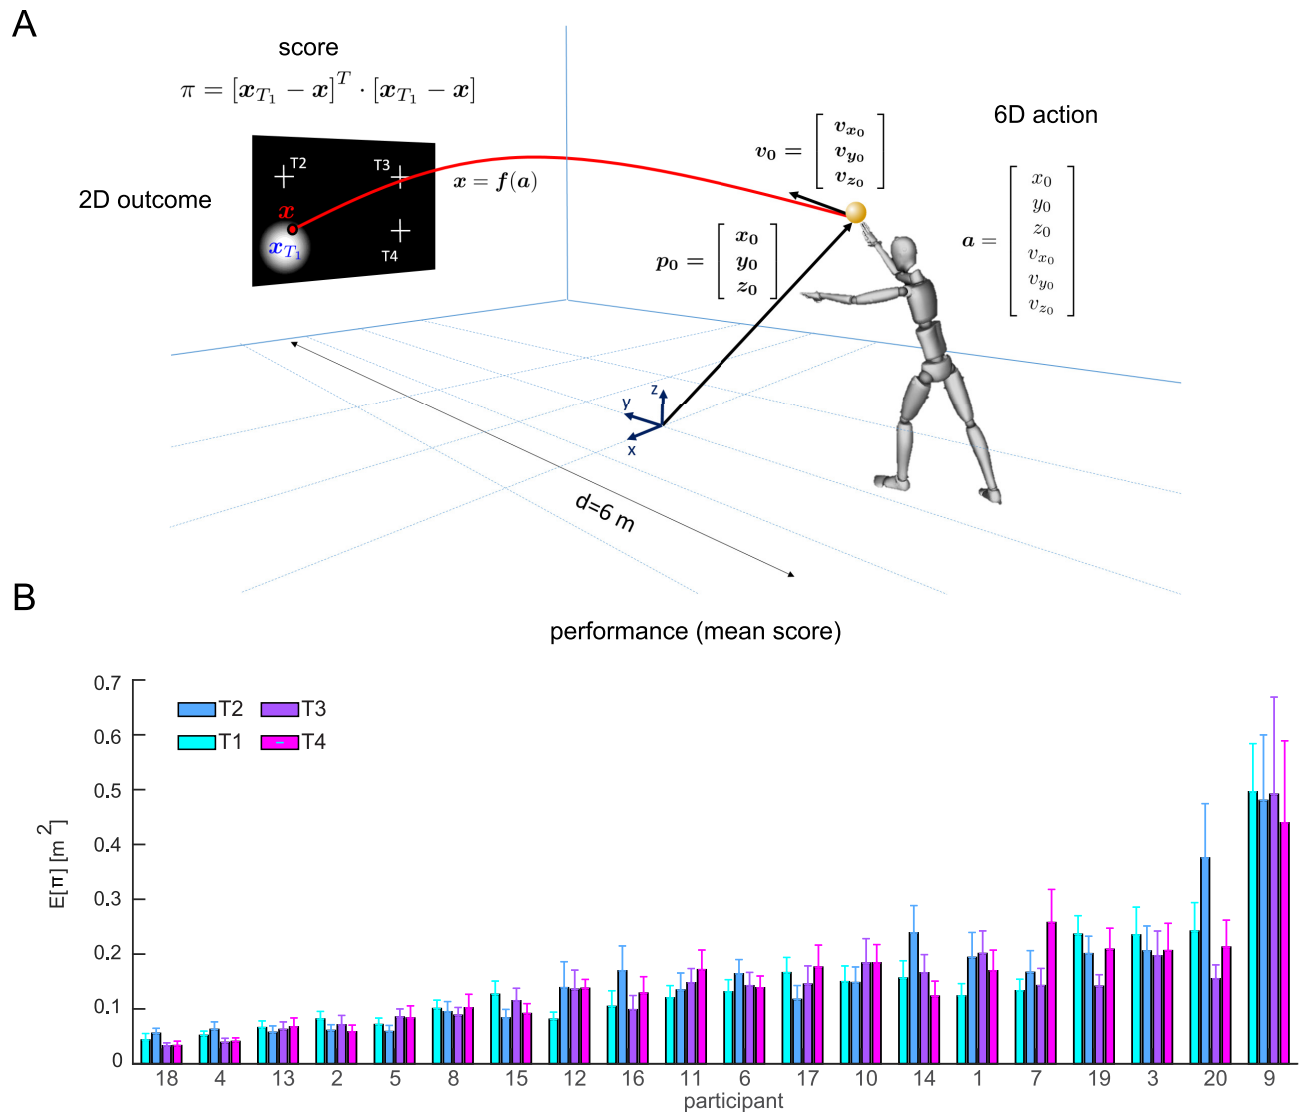


Fig 1. Unconstrained throwing task and experimental performance. (A) Schematic representation of the unconstrained overarm throwing task considered in this study as an example of a complex motor skill. The release parameters characterizing a throwing action are six-dimensional (three position components and three velocity components). (B) Performance (mean squared error \pm SE) across participants ($n = 20$) and four targets.

<https://doi.org/10.1371/journal.pone.0253626.g001>

actions, i.e. factors depending on the local geometry of the goal function, rather than only on different amounts of action variability.

The impact of the interplay between motor variability and task geometry on performance in goal-directed behaviors has also been investigated with an approach that does not require any assumption on the smoothness of the action-to-score mapping, but relies on the computation of surrogate data or on an exhaustive search of optimal data distributions, rather than analytic descriptions, to explore its geometry [7, 11, 20]. The Tolerance-Noise-Covariation (TNC) method [7] allows to quantify the difference in performance between two series of trials as the sum of three contributions. The *tolerance* component is associated with potential differences in the *sensitivity* of the local action-to-score mapping geometry associated with different choices of the mean action. The *noise* component quantifies the impact of different amount of action variability in the two series. Finally, the *covariation* component accounts for the impact

of different alignment of the action variability with the local geometry. A revised version of the method (TNC-Cost analysis [20]) allows to assess the three components of a single series of trials with respect to optimal performance.

A key aspect of the TNC approach that makes it particularly suitable for the analysis of inter-individual differences in goal-directed behaviors is its focus on the relationship between action distribution and performance as mean score. While the UCM and GEM approaches focus on variability in action space and outcome space, the TNC decomposition shows how the mean score depends on the choice of the mean action in addition to variability. Furthermore, by identifying different contributions to the mean score, the TNC approach allows to characterize individual performances with a higher level of details: for instance, a participant could perform the task with a higher variability than a peer, but achieve the same level of performance thanks to a better alignment. However, because the computations of the costs depend on numerical procedures that become cumbersome for high dimensional action spaces [11], the TNC analysis has been applied only to simple tasks, such as a planar virtual version of the “skittles” game in which one can control the angle and velocity of ball release [7, 21].

In the present work, our goal was to characterize inter-individual differences in the relationship between action-variability and performance in a complex, real-life motor skill. We asked twenty non-expert participants to perform unconstrained overarm throws at four different targets placed on a vertical plane at 6 m distance, as depicted in Fig 1A. When we analyzed the time-course of the whole-body kinematics during the throwing motion [22, 23], we found that throwing styles, i.e. the average trajectory of each limb segment, differed considerably across individuals. Similarly, we found large differences in individual performances, as shown in Fig 1B. Here, we focused on inter-individual differences in terms of release parameters distribution and throwing performance, as differences in throwing styles may translate into different release strategies, which may or may not correspond to differences in the mean score. We aimed at characterizing the key performance-related features of the individual release parameters distribution. We also aimed at identifying consistent features that could explain why participants differed in performance and how they could achieve similar performance levels with different strategies.

To identify the different contributions of the distribution of throwing actions to performance, we introduced a novel analytic method that can be applied to an unconstrained throwing task, described by at least six release parameters, overcoming the computational limitation of the TNC method, which requires a number of numerical operations that scale exponentially with the number of dimensions of the action space. Our approach is based on second order derivatives of the action-to-score function and depends on the following two assumptions: i) the action distribution is sufficiently localized in a region of the action space; ii) the score function, although non-linear, can be adequately approximated with a second-order Taylor-expansion. Hence, we make use of the Hessian matrix, i.e., the matrix of second-order partial derivatives, rather than the Jacobian, to estimate the local tolerance of the score as well as to estimate the alignment of action covariance with the curvature of the action-to-score mapping.

2 Results

To relate performance in a goal-directed motor task to the distributions of motor actions, we introduce a decomposition of the mean score in terms of a few parameters depending on the mean action, the Hessian of the action-to-score mapping computed at the mean action, and the covariance of the action distribution. Fig 2A illustrates the variables and functions describing the relationship between an action (\mathbf{a}) and a score (or loss, π) assigned to the outcome (\mathbf{x}) of the action. Such score represents a scalar measure of inaccuracy of the outcome with respect

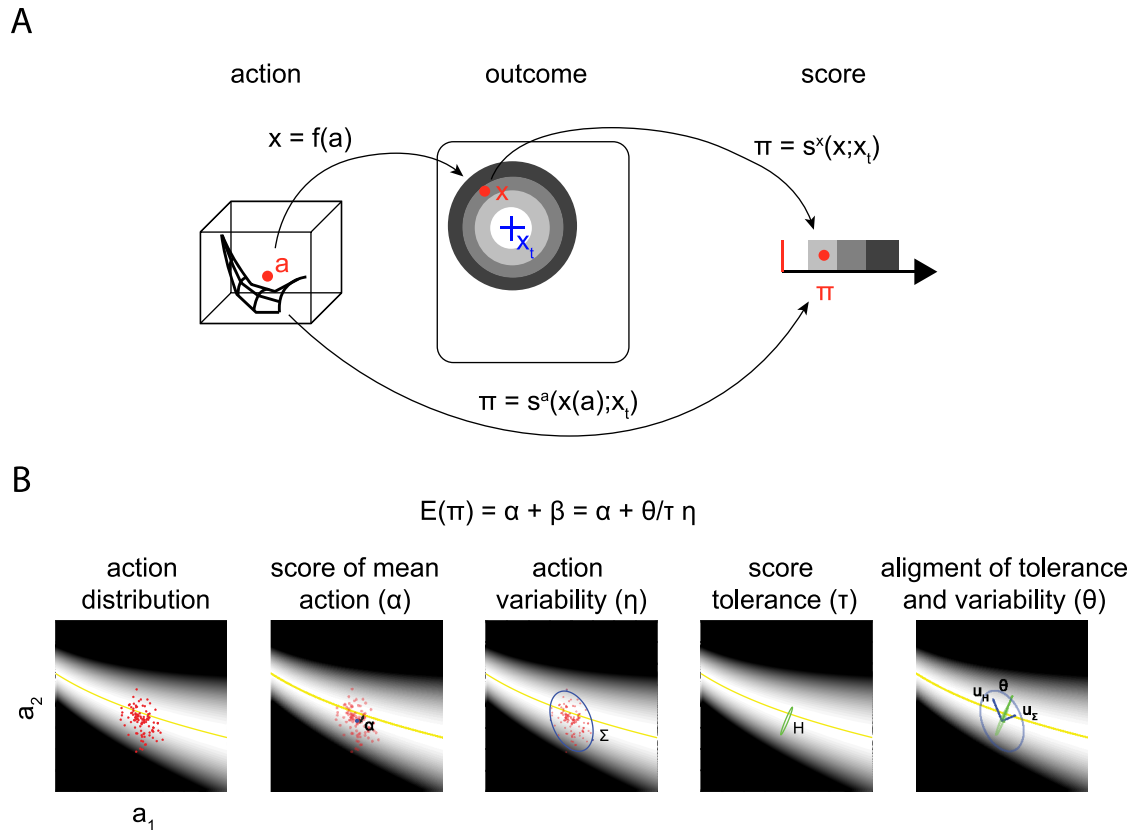


Fig 2. Definition of variables (action, outcome, and score) and schematic overview of the proposed approach to decompose the mean score (performance). (A) An action, represented by a vector \mathbf{a} in a high-dimensional action space (illustrated by a red circular marker in a three-dimensional action space), leads to an outcome, $\mathbf{x} = f(\mathbf{a})$, i.e. a vector typically represented in outcome or goal space with less dimensions. The outcome is then associated with a corresponding score, $\pi = s^x(\mathbf{x}; \mathbf{x}_T)$, based on its relation with respect to the aimed target. The set of points in action space associated with an outcome achieving the goal constitute the solution manifold (illustrated by the black mesh). (B) Illustration of the Hessian-based decomposition of the mean score in the case of two-dimensional actions. The action-to-score mapping is represented by the gray-shaded areas. A set of actions (small red markers, first panel) has a distribution characterized by the mean (large blue marker, second panel) and covariance (Σ , blue ellipse in the third and fifth panels). The tolerance of the score to variations of the actions around their means is characterized by the Hessian (H , green ellipse in the fourth panel). The mean score ($E(\pi)$) is decomposed as the sum of the score of the mean action (α) and a tolerance-variability index (β) expressed as the product of three terms: the reciprocal of tolerance (τ), noise (η), and alignment (θ), i.e. a scalar characterizing whether the most sensitive directions of the Hessian (green arrows \mathbf{u}_H , fifth panel) are aligned to the directions of highest action variability (blue arrows \mathbf{u}_Σ).

<https://doi.org/10.1371/journal.pone.0253626.g002>

to the goal (or target, \mathbf{x}_T) and, thus, it is a composite function of the action-to-outcome mapping ($\mathbf{x} = f(\mathbf{a})$) and the outcome-to-score mapping ($\pi = s^a(\mathbf{x}; \mathbf{x}_T)$).

If we consider an action distribution with mean $\bar{\mathbf{a}}$ and covariance Σ^a , by using a second-order Taylor-expansion of the action-to-score function around the mean action (see Methods), the mean score can be approximated as:

$$\bar{\pi} \approx \alpha + \beta \tag{1}$$

where

$$\alpha = \alpha(\bar{\mathbf{a}}) \equiv s^a(\bar{\mathbf{a}}) \tag{2}$$

is the score of the average action, and

$$\beta = \beta(\bar{\mathbf{a}}, \Sigma^a) \equiv \text{trace}\left(\frac{1}{2}H(\bar{\mathbf{a}})\Sigma^a\right) \tag{3}$$

is a *tolerance-variability index* which captures how local non-linearities of the action score (the Hessian $H(\bar{\mathbf{a}})$) and variability in the action strategy (the covariance Σ^a) affect performance.

Furthermore, by defining *score sensitivity* as the trace of H and *score tolerance* as its reciprocal

$$\tau = \tau(\bar{\mathbf{a}}) \equiv \frac{1}{\text{trace}(H(\bar{\mathbf{a}}))} \tag{4}$$

and *amount of noise* or, briefly, *noise* as the trace of the covariance matrix (i.e., the total variation of the action distribution, providing a scalar measure of the overall amount of action variability)

$$\eta = \eta(\Sigma^a) \equiv \text{trace}(\Sigma^a), \tag{5}$$

we can express the *tolerance-variability index* as

$$\beta = \text{trace}\left(\frac{1}{2}\bar{H}\bar{\Sigma}\right)\frac{\eta}{\tau} = \frac{\theta}{\tau}\eta \tag{6}$$

where we introduced the *alignment* as

$$\theta = \theta(\bar{\mathbf{a}}, \Sigma^a) = \text{trace}\left(\frac{1}{2}\bar{H}\bar{\Sigma}\right) \tag{7}$$

with $\bar{H} = H/\text{trace}(H)$ and $\bar{\Sigma} = \Sigma/\text{trace}(\Sigma)$. The alignment is here a scalar that measures the effect on performance of the relative orientation between the directions of highest sensitivity of the action-to-score mapping with respect to the directions of highest variability of the actions (see [Methods](#)).

[Fig 2B](#) provides a graphical illustration of the decomposition approach in the case of two-dimensional actions, as in a throwing task in which only the horizontal and vertical components of the release velocity can vary. The score associated to each action is indicated by the *gray-shaded* background. The yellow curve at the center of the white area represents the *solution manifold*, the set of actions that accurately achieve the goal. The distribution of actions is characterized by their mean ($\bar{\mathbf{a}}$) and covariance (Σ) while the local geometry of the action-to-score mapping is characterized by the Hessian (H) computed at the mean action.

In sum, the expected value of the score ($E(\pi)$) can be decomposed as the sum of the score of the mean action (α) and the tolerance-variability index (β)

$$E[\pi] \approx \alpha + \beta = \alpha(\bar{\mathbf{a}}) + \frac{\theta(\bar{\mathbf{a}}, \Sigma^a)}{\tau(\bar{\mathbf{a}})}\eta(\Sigma^a) \tag{8}$$

where β is determined by the interplay between the sensitivity of the performance to action variability (reciprocal of the tolerance τ), the magnitude of the action variability η , the alignment between the Hessian and the action covariance θ .

2.1 Simulated 2D throwing examples

To illustrate how our decomposition allows to identify the key features of the action distribution contributing to performance, we consider a toy model of a throwing task in which a

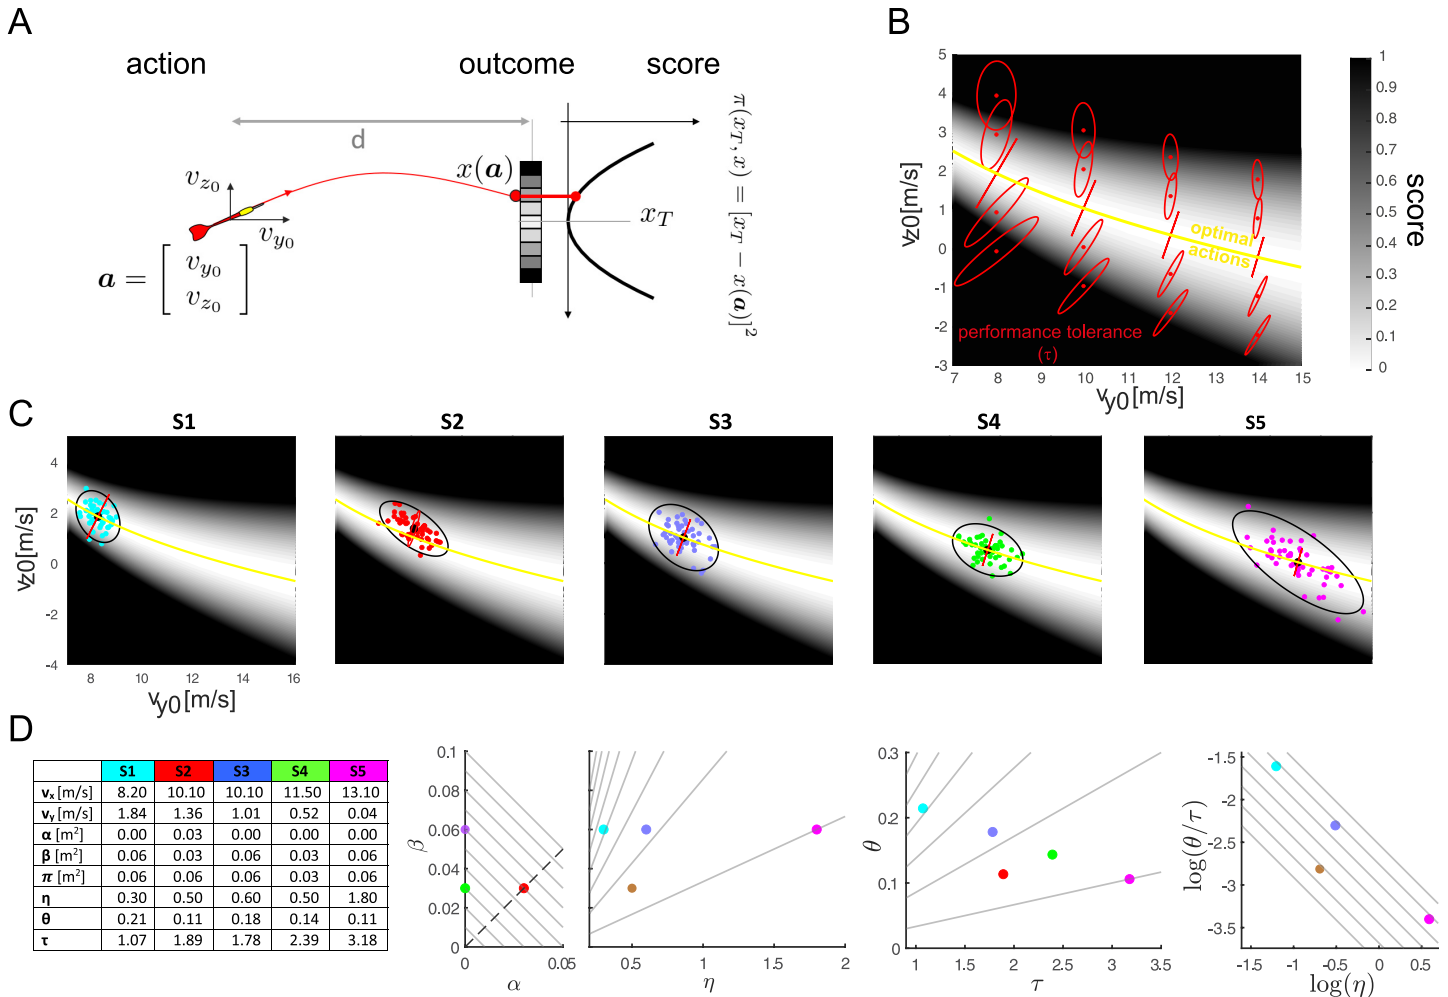


Fig 3. Examples of different throwing strategies in a simulated 2D throwing task. (A) Toy model of a throwing task in which the action (\mathbf{a}) is characterized by only two parameters, i.e. the horizontal (v_{y0}) and the vertical (v_{z0}) release velocity components, the outcome (x) is the arrival position of the projectile on a vertical plane at a distance d from the release position, and the score (π) is the squared distance of the arrival position from a target (x_T). (B) Action score (gray-shaded background), solution manifold (yellow line), and Hessian (red ellipses) in a region of release parameters. The score sensitivity for different throwing strategies is captured by the Hessian, which is represented with red ellipses whose major axis indicates, locally, the direction of maximum sensitivity (smaller tolerance to noise). (C) Five (simulated) individual strategies. For each strategy (S1-S5, different panels and colors) the individual throwing actions are shown together with mean (large black circle), their covariance (black ellipse, 95% C.L.) and the Hessian at the mean action (red ellipse). (D) Hessian-based decomposition of the mean score of the five strategies. For each strategy, the mean release parameters (v_{y0} and v_{z0}), the mean score ($\bar{\pi}$), and the indicators from the decomposition (score of the mean action α , tolerance-variability index β , tolerance (τ), noise η , and alignment θ) are presented in the table. The four scatter plots illustrate how these indicators allow to differentiate strategies with different performance as well as strategies with same performance. See text for more details.

<https://doi.org/10.1371/journal.pone.0253626.g003>

projectile is released from a fixed position with given horizontal and vertical velocity components and must hit a target at a fixed height on a vertical plane at a fixed distance. Additionally, we assume that the horizontal component of the release velocity lies in the vertical plane containing the release position and the target, so that the action \mathbf{a} is a two-dimensional (2D) vector, the projectile trajectory lies on a vertical plane, and the outcome (x , i.e., the arrival height) is a scalar (Fig 3A). The score $\pi(x; x_T)$ is the squared distance of the arrival position from the target.

As in Fig 2B, since the actions are 2D, in Fig 3B we can represent the action-to-score function as contours in the action plane (i.e., the plane v_{y0} - v_{z0} of the horizontal and vertical release velocities) with gray shading (lighter shades for lower error) and the solution manifold, i.e. the

set of optimal actions which results in 0 penalty score, with a yellow curve. Fig 3B also shows the local sensitivity of the score, i.e. the key element of our decomposition, which is captured by the Hessian of the action-to-score function. For different values of the mean action (i.e. \bar{v}_{y_0} and \bar{v}_{z_0}), the sensitivity of the mean score is represented by a red ellipse. The longer axis of the ellipse, whose length is proportional to the largest eigenvalues of the Hessian matrix (see Methods), corresponds to the direction in which a given deviation of an action from the mean affects the mean score the most, i.e. the direction of maximal sensitivity or minimal tolerance to noise. The length of each axes of the ellipse indicates how much a deviation in that direction affects the mean score. Consequently, action distributions with the same amount of noise will have different performances depending on the alignment of their covariance with the direction of maximal sensitivity and on the lengths of the axes of the local sensitivity ellipses. Note that the orientation of the ellipses, their size, and the ratio of their axes depend on the position on the action plane. On the solution manifold, the ellipses have only one axis with non-zero length, indicating that there is only one direction in action space which is relevant for the score. This direction however, is not constant, but changes along the solution manifold. Also note that the ellipses are smaller for larger horizontal velocities v_{y_0} , indicating that the score is more tolerant for faster throws.

2.1.1 Iso-performing strategies: Trade-offs between indicators derived from the Hessian-based decomposition. Fig 3C shows simulated action distributions for five different strategies, each represented by 50 throws with v_{y_0} and v_{z_0} drawn from a bi-dimensional Gaussian distribution with a different mean and a different covariance. The plots are arranged from left to right according to the magnitude of the mean horizontal velocity (v_{y_0}). The fourth strategy (S4 green) is the best performing of the five, as it has a mean action on the solution manifold and the direction of maximal action variability (major axis of the black covariance ellipse) is almost orthogonal to the direction of maximal noise sensitivity (major axis of the red local sensitivity ellipse). The second strategy (S2 red) has a worse performance than the fourth, even if the two have the same magnitude of noise η and a similar low alignment of the covariance with the local noise sensitivity. This is because its mean action is not on the solution manifold, as it has a vertical release velocity larger than optimal. The remaining three strategies (S1 cyan, S3 blue, S5 magenta) have different amount of noise and different levels of alignment with the local noise sensitivity but they have the same performance, also equal to that of S2. However, it is difficult to recognize by visual inspection of the action distribution that the four strategies (S1, S2, S3, and S5) have the same performance and to characterize their differences.

The Hessian-based decomposition of the mean score of the five strategies according to Eqs (1) and (6) is illustrated in Fig 3D. Four scatter plots between pairs of the different indicators derived from the decomposition (score of the mean action α , tolerance-variability index β , tolerance τ , noise η , and alignment θ) allow to compactly describe the five strategies, to directly relate them to performance, and to characterize their differences. The α - β scatter plot shows that strategies with the same performance may trade-off accuracy of the mean action (α) with the combined effect of action variability and local geometry of the action-to-score mapping (β). The diagonal lines with slope -1 represent *iso-performing strategies* (constant $\alpha + \beta$ equal to different values of $\bar{\pi}$). Strategy S4 is the best performing one as it has zero α and a low value of β . Strategy S2 has the same low β as S4 but a higher $\bar{\pi}$ due to non-zero α . S2 performance matches that of the remaining three strategies (S1, S3, and S5) which have higher β but null α . In fact, the four strategies lie on the same diagonal line.

The η - β scatter plot reveals that S1, S3, and S5 achieve the same β with different amounts of noise (η). This occurs because their action distributions differ in the alignment of the direction of largest variability with the most noise sensitive directions (θ) and because their mean action

is located in regions of the action plane with different noise tolerance (τ), as shown in the τ - θ scatter plot. The lines in the η - β scatter plot represent strategies that have the same slope in the linear dependence of the tolerance-variability index β on noise η , because they have a given ratio between alignment θ and tolerance τ , as $\beta = \frac{\theta}{\tau}\eta$. Analogously, the lines in the τ - θ scatter plot represent strategies with equal $\frac{\theta}{\tau} = \frac{\beta}{\eta}$. As the horizontal speed increases, the distributions are located in regions with higher tolerance τ and they have a lower alignment θ .

The trade-off between the effect of variability and the effect of local score sensitivity can be best illustrated in the $\log(\eta)$ - $\log(\frac{\beta}{\tau})$ scatter plot. Since $\log(\beta) = \log(\frac{\theta}{\tau}) + \log(\eta)$, the diagonal lines with slope -1 represent strategies in which the same β results from an increase in noise η compensated by a decrease in alignment θ or an increase in tolerance τ . S1 (*cyan*) is in a region with lower tolerance ($\tau = 1.07$, see table in Fig 3D) and has a higher alignment ($\theta = 0.21$), but it compensates for such unfavorable conditions with lower noise ($\eta = 0.30$). S3 (*blue*) has a larger mean horizontal release and a corresponding higher tolerance ($\tau = 1.78$) and it achieves the same performance with larger noise ($\eta = 0.60$). Finally, S5 (*magenta*) tolerates even larger noise ($\eta = 1.80$) by virtue of the higher tolerance ($\tau = 3.18$) associated to the largest mean release horizontal velocity.

In sum, the four iso-performing strategies (S1, S2, S3, and S5) illustrate two different trade-offs between indicators derived from the Hessian based decomposition. S2 vs. S1, S3, and S4 trade-off bias with tolerance-variability. S1 vs. S3 vs. S5 trade-off noise alignment and sensitivity with noise magnitude. In both cases, the indicators allow to compactly characterize the features of the action distribution determining a given performance and those distinguishing different strategies, even when they have the same performance.

2.1.2 Comparison with TNC, UCM, and GEM analyses. The 2D throwing toy model also allows to compare the Hessian-based decomposition with existing computational methods that have been introduced for analyzing the relationship between an action distribution and its mean score (TNC), for decomposing action variability into task-relevant and task-irrelevant components (UCM), and for relating action variability and outcome variability (GEM).

The TNC-Cost analysis [20] characterizes an action distribution with three indicators (or costs) computed by comparing the mean score of the distribution with the score of a transformed distribution, optimal according to different criteria. The *tolerance* cost (T-cost) represents the difference between the mean score of the distribution and the mean score of the same distribution translated to an optimal location, i.e. a location with the lowest possible score. For the *noise* cost (N-cost) the optimal distribution is obtained by shrinking the distribution (i.e., reducing the noise η) without affecting the mean. The *covariation* cost (C-cost) is computed using an optimal distribution with the same mean and the same individual action components (e.g. v_{y_0} and v_{z_0} in our 2D throwing example) but with an optimal re-association of those components in different actions, i.e. with an optimal covariation. These indicators are related to the indicators of our decomposition (see Section A.2 in S1 Appendix) but they represent a re-combinations of those indicators and, more importantly, do not provide a direct decomposition of the mean score, thus they do not allow to describe the performance trade-offs between indicators presented above. Also, critically, as the computation of both the N-Cost and the C-cost relies on exhaustive searches, they require a number of operations that scales exponentially with the dimensions of the action space and they are not feasible for 6D actions, as in unconstrained throwing.

Fig 4A presents an example of three simulated strategies with the same performance and an α - β trade-off. Increasing the mean horizontal release velocity v_{y_0} , a larger bias α is compensated by a lower tolerance-variability index β achieved by a lower alignment θ and higher tolerance τ with constant noise η . The T-Cost is constant, as the three distributions have the same

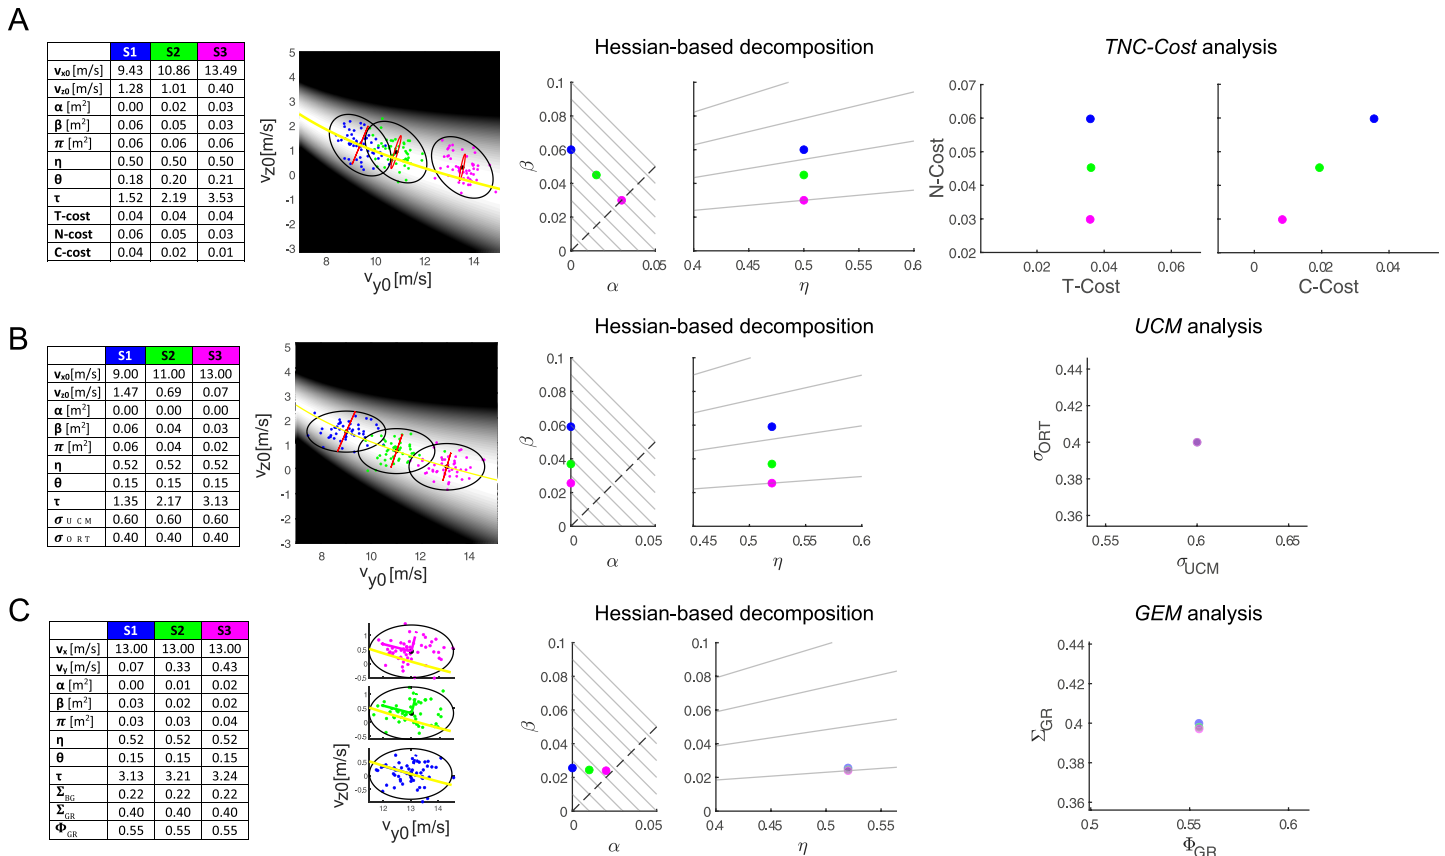


Fig 4. Comparison of the Hessian-based decomposition with existing computational methods for the analysis of throwing strategies in a 2D throwing task. (A) Comparison with TNC-Cost analysis. Three strategies (S1 blue, S2 green, S3 magenta) are simulated by generating a distribution of 50 release parameters (v_{y0} , v_{z0}) by randomly sampling from a bi-dimensional Gaussian distribution with a given mean and covariance (see parameters in the table on the left). The strategies have the same performance because they trade-off bias (α) with tolerance-variability (β), as can be clearly observed in the α - β scatter plot. Moreover, the strategies have equal noise (η). However, the strategies have equal T-cost but different N-Cost, which then represent a recombination of the indicators extracted from the Hessian-based decomposition. Moreover, as the TNC costs do not sum up to the performance, the α - β tradeoff cannot be observed. (B) Comparison with UCM analysis. Three simulated strategies with no bias (i.e. with a mean action on the solution manifold) differ in performance because of differences in β , due to differences in the alignment of the covariance (black ellipses) with the Hessian (red degenerate ellipses), since the noise is equal. The decomposition of the action variability onto an *uncontrolled manifold* (UCM) subspace and an orthogonal (ORT) subspace does not discriminate the three strategies. (C) Comparison with GEM analysis. Three simulated strategies with equal noise but different performance due to differences in α are not discriminated by the *goal-relevant sensitivity* Σ_{GR} and the *goal-relevant variability fraction* Φ_{GR} .

<https://doi.org/10.1371/journal.pone.0253626.g004>

mean score and the same covariance matrix, i.e. they differ only for the position of the mean action, and therefore they have the same difference with the mean score of the optimally translated distribution. Thus, the T-Cost does not distinguish iso-performing strategies with different bias when this is compensated by better alignment and tolerance. Moreover, the strategies have different N-cost, even if their action distribution have the same covariance matrix. This is due to the fact that when the action-to-score function is smooth and the dispersion of the actions not too high, as in this case, the optimal distribution used to compute the N-cost has zero noise, i.e. is concentrated on the mean action, and the N-cost corresponds to $\beta = \bar{\pi} - \alpha$.

In contrast to both Hessian-based decomposition and TNC analysis, UCM analysis [15, 16] aims at decomposing action variability rather than characterizing the dependence of mean score on strategy and variability. It decomposes variability into a component within an *uncontrolled manifold* (UCM) subspace, which does not affect the task because such subspace is the null space of the Jacobian of the action-to-outcome function, and a component within the orthogonal (ORT) subspace. The ratio $\frac{\sigma_{UCM}}{\sigma_{ORT}}$, where σ_{UCM} is the square root of the variance in

the UCM and σ_{ORT} in the ORT, then indicates whether a controller is successful in pushing variability onto task-irrelevant dimensions, thus improving performance. However, as shown by our decomposition, the relationship between action variability and performance, quantified by mean score, depends on the alignment of the action covariance with the Hessian and on the size of the Hessian rather than on the alignment with the Jacobian. Thus, a decomposition of variability onto the UCM and ORT subspaces does not discriminate strategies with different performance but equal σ_{UCM} and σ_{ORT} (and therefore equal ratio).

Fig 4B shows an example of three simulated strategies with different mean horizontal release velocity and different performances, due to different β , but equal σ_{UCM} and σ_{ORT} . The three action distributions have means on the solution manifold, i.e. $\alpha = 0$ and equal noise η . They also have the same orientation (with major and minor axes parallel to the v_{y_0} and v_{z_0} axes respectively) of the covariance ellipses but different ratios between the axes of the ellipses. As the Jacobian as well as the associated UCM and ORT directions change with the mean release velocity, the different amounts of variability along the horizontal and vertical velocities have the same amount of variability projected along the UCM and ORT directions. In contrast, the changes in the alignment of the covariance with the Hessian, which also changes with the mean release velocity, as well the changes in the size of the Hessian, do capture the features of the action distribution that affect performance. In particular, the differences in mean score, are due to different ratios of θ and τ as shown by the iso-lines in the $\beta - \eta$ plane.

Finally, the GEM analysis [17] extends the UCM analysis by considering the noise sensitivity of the action-to-outcome function in addition to the directions of the Jacobian and its null space. The *total body-goal sensitivity* at the mean action, i.e., the ratio $\Sigma_{BG} = \frac{\sigma_e}{\sigma_u}$, where σ_e is the square root of the total variation of the outcome ($\sigma_e = \sqrt{\text{trace}(\Sigma_x)}$ in our notation) and σ_u is the square root of the total variation of the action ($\sigma_u = \sqrt{\eta}$), gives the amplification of variability between the action and outcome levels and can be decomposed as

$$\Sigma_{BG} = \Sigma_{GR} \Phi_{GR} \tag{9}$$

where the *goal-relevant sensitivity* $\Sigma_{GR} = \frac{\sigma_e}{\sigma_R}$, with $\sigma_R = \sigma_{ORT}$, gives the amplification of variability in the ORT subspace and the *goal-relevant variability fraction* $\Phi_{GR} = \frac{\sigma_R}{\sigma_u}$ is the fraction of action variability available for mapping into the outcome and hence that can have an effect on performance. However, as for the UCM analysis, the GEM analysis characterizes the effect of the first order approximation of the action-to-outcome function, i.e., the Jacobian, on the variability in outcome space rather than the effect of the interplay between the geometry of the action-to-score function and the action distribution (both mean and covariance) on the mean score, i.e. the performance. Indeed, the GEM does not capture the effect on performance of a non-optimal mean in the distribution, which gives rise to the α - β trade-off presented above.

Fig 4C shows an example of three simulated strategies with the same horizontal release velocity and the same noise but different mean vertical release velocities and, thus, different α and different performances, as clearly shown in the α - β scatter plot. In contrast, the GEM decomposition does not discriminate between the three strategies, as they all have nearly equal Φ_{GR} and Σ_{GR} . This is due to the fact that covariances are also nearly equal and the Jacobian does not change much when the mean vertical release velocity deviates from the optimal value on the solution manifold. However, such deviation of the mean action has a substantial effect on performance.




Fig 5. Examples of distribution of release velocity components for five representative participants (target T1). Each row illustrates a pairs of release velocity components ($v_{x_0} - v_{y_0}$, $v_{x_0} - v_{z_0}$, $v_{y_0} - v_{z_0}$). Colored circles represent release parameters of individual throws. Black circles and ellipses represent mean and covariance (95% c.l.) of each parameter distribution. The gray-level shading of the contours indicates the local score, as a function of the release parameters. Note that in each row, the range of horizontal and vertical axes are the same for all participants, and the different shapes of the score across participant, reflect individual differences in the average action (mean position and velocity vectors at release). The wider the white area around the mean action, the more tolerant the score is to noise. Participants have been sorted from left to right according to their average release speed.

<https://doi.org/10.1371/journal.pone.0253626.g005>

2.2 Unconstrained 6D throwing experimental data

To gain insights on individual strategies and their relationship to performance in a complex real-world motor skill we applied our Hessian-based decomposition to unconstrained throwing. We instructed 20 participants to perform unconstrained overarm throws to 4 circular targets arranged on a vertical target board placed at 6 m from the initial position (see [Methods](#)). We considered the release position and velocities ([Fig 1A](#)) and the squared distance of the arrival position from the center of the target as the score (loss) of each trial. As we found that performance varied significantly across participants ([Fig 1B](#), $p < 0.001$ linear mixed-effect model with participant and target identity as categorical predictors), our aim was to characterize the key features of the individual action distributions that relate to performance and to identify features that differentiate participants without affecting performance.

As actions are six-dimensional (6D), it is impossible to visualize both the score and the individual strategies in a single plot, as in the 2D throwing case. To illustrate some of the characteristics of the individual action distributions, [Fig 5](#) shows the distribution of the three pairs of release velocity components in five exemplar participants (*columns*) aiming at target T1. As there are 15 different pairs of 6 action variables (3 position and 3 velocity components) these plots do not fully illustrate the individual strategies (see Section B in [S1 Appendix](#) for a display

of 9 of the 15 possible pairs). However, these examples show that the action distributions and the local geometry of the score function differ significantly across participants even when we consider a subset of action variables.

As for the simulated 2D strategies illustrated in Fig 3C, participants have been sorted from left to right according to their mean longitudinal release velocity (v_{y_0}), hence P10 is the slowest thrower and P11 the fastest. The domain of each velocity variable corresponds to the population mean ± 3 standard deviations. The score associated to each pair of velocity components, with the remaining velocity component and the three position components fixed at the subject-specific mean value, is indicated by the gray-scale shading of the contours. Note that the regions corresponding to actions with the same score (i.e. same gray shading) have different shapes (or geometry) across planes and participants, as they depend on the individual mean release action. The wider the white areas around the mean release parameter (*black large circle*), the more tolerant is the action-to-score function to action variability. In the $v_{y_0} - v_{z_0}$ plane (*third row*), as in the 2D case, tolerance increases with mean v_{y_0} . However, the mean longitudinal release velocity also affects the tolerance along the horizontal transverse release velocity (v_{x_0}), as can be observed in the $v_{x_0} - v_{y_0}$ (*first row*) and $v_{x_0} - v_{z_0}$ (*second row*) planes.

The distributions of the release parameters (*colored circles*), summarized in each plot of Fig 5 in terms of mean (*black large circle*) and covariance (two-standard deviations, *black ellipse*), also differs across planes and participants. In addition to differences in mean longitudinal velocity (v_{y_0}), there are differences in mean vertical velocity (v_{z_0}). In particular, P15 (*second column*), on average, releases the ball with a faster vertical velocity than P1, who has a similar longitudinal velocity. Remarkably, individual covariances differ both in size and orientation. In the $v_{y_0} - v_{z_0}$ (*third row*) plane, P10 shows a small covariance ellipse with the major axis oriented along the v_{z_0} axis. In contrast, P11 shows a much larger covariance ellipse with the major axis oriented diagonally, i.e. with a negative correlation between the two velocity components.

In sum, in all three pairs of velocity variables (as well as in the other pairs of variable shown in Section B in S1 Appendix), both the local geometry of the action-to-score function and the structure of action variability suggest that specific features of the individual action distributions affect performance, as they determine the amount of overlap of the distribution with the regions with lowest penalty. However, as the distribution is 6-dimensional and there are 15 different pairs of action variables, it is not possible to identify by visual inspection a unique source of the inter-individual differences in throwing strategies and to systematically explain the relationship between action distribution and throwing performance. These limitations can be overcome by the Hessian-based decomposition.

2.2.1 Validity of the Hessian-based decomposition. The Hessian-based performance analysis is based on the assumption that the individual action distribution is sufficiently localized around the mean action such that higher order terms of the Taylor expansion can be neglected in the action score approximation. To validate this assumption with our experimental data, we tested whether Eq (8) can be considered as an acceptable model of the mean action score. Fig 6 shows the relationship between the mean squared error $\bar{\pi}$ and the sum $\alpha + \beta$ across participants and targets. We found that the sum $\alpha + \beta$ could explain 99% of the variance for targets T2, T3, and T4 and 97% for target T1, due to the larger error observed in P9. P9 had in fact the worst performance (see Fig 1B) and the largest bias (α) and variability (η), which violates the assumption of the action distribution being sufficiently localized. In sum, $\alpha + \beta$ can explain quite accurately the individual performance across participants and conditions, except for P9, where the sum $\alpha + \beta$ tends to overestimate the average score for T1, T2 and T3.

2.2.2 Hessian-based decomposition of the mean score reveals key performance-related features of individual throwing strategies. The Hessian-based decomposition of the mean




Fig 6. Validity of the quadratic approximation. Mean score vs. quadratic approximation (8) across targets (T1-T4, different panels) and participants (P1-P20) as in Fig 1. Note that P9 is the only participant deviating substantially from the identity line (dashed line).

<https://doi.org/10.1371/journal.pone.0253626.g006>

score relies on the straightforward computation of the action mean, action covariance and Hessian of the action-to-score function at the mean action. While the analysis of the structure of the Hessian and the action covariance matrices provides useful information on the task geometry and on the features of the action variability shared across subjects (see Section C in [S1 Appendix](#)), the parameters extracted from the Hessian-based decomposition of the mean score compactly describe the key performance-related features of the individual action distribution and allow to fully characterize how performance depends on such features.

[Fig 7](#) shows the distributions of all the parameters of the Hessian-based decomposition of the mean score (α , β , η , τ and θ defined in [Eq 8](#)) across participants and relative to target T1. To visualize the inter-individual differences in terms of these five parameters, as for 2D throwing, we consider four scatter plots. The five exemplar participants of [Fig 5](#) are identified by the

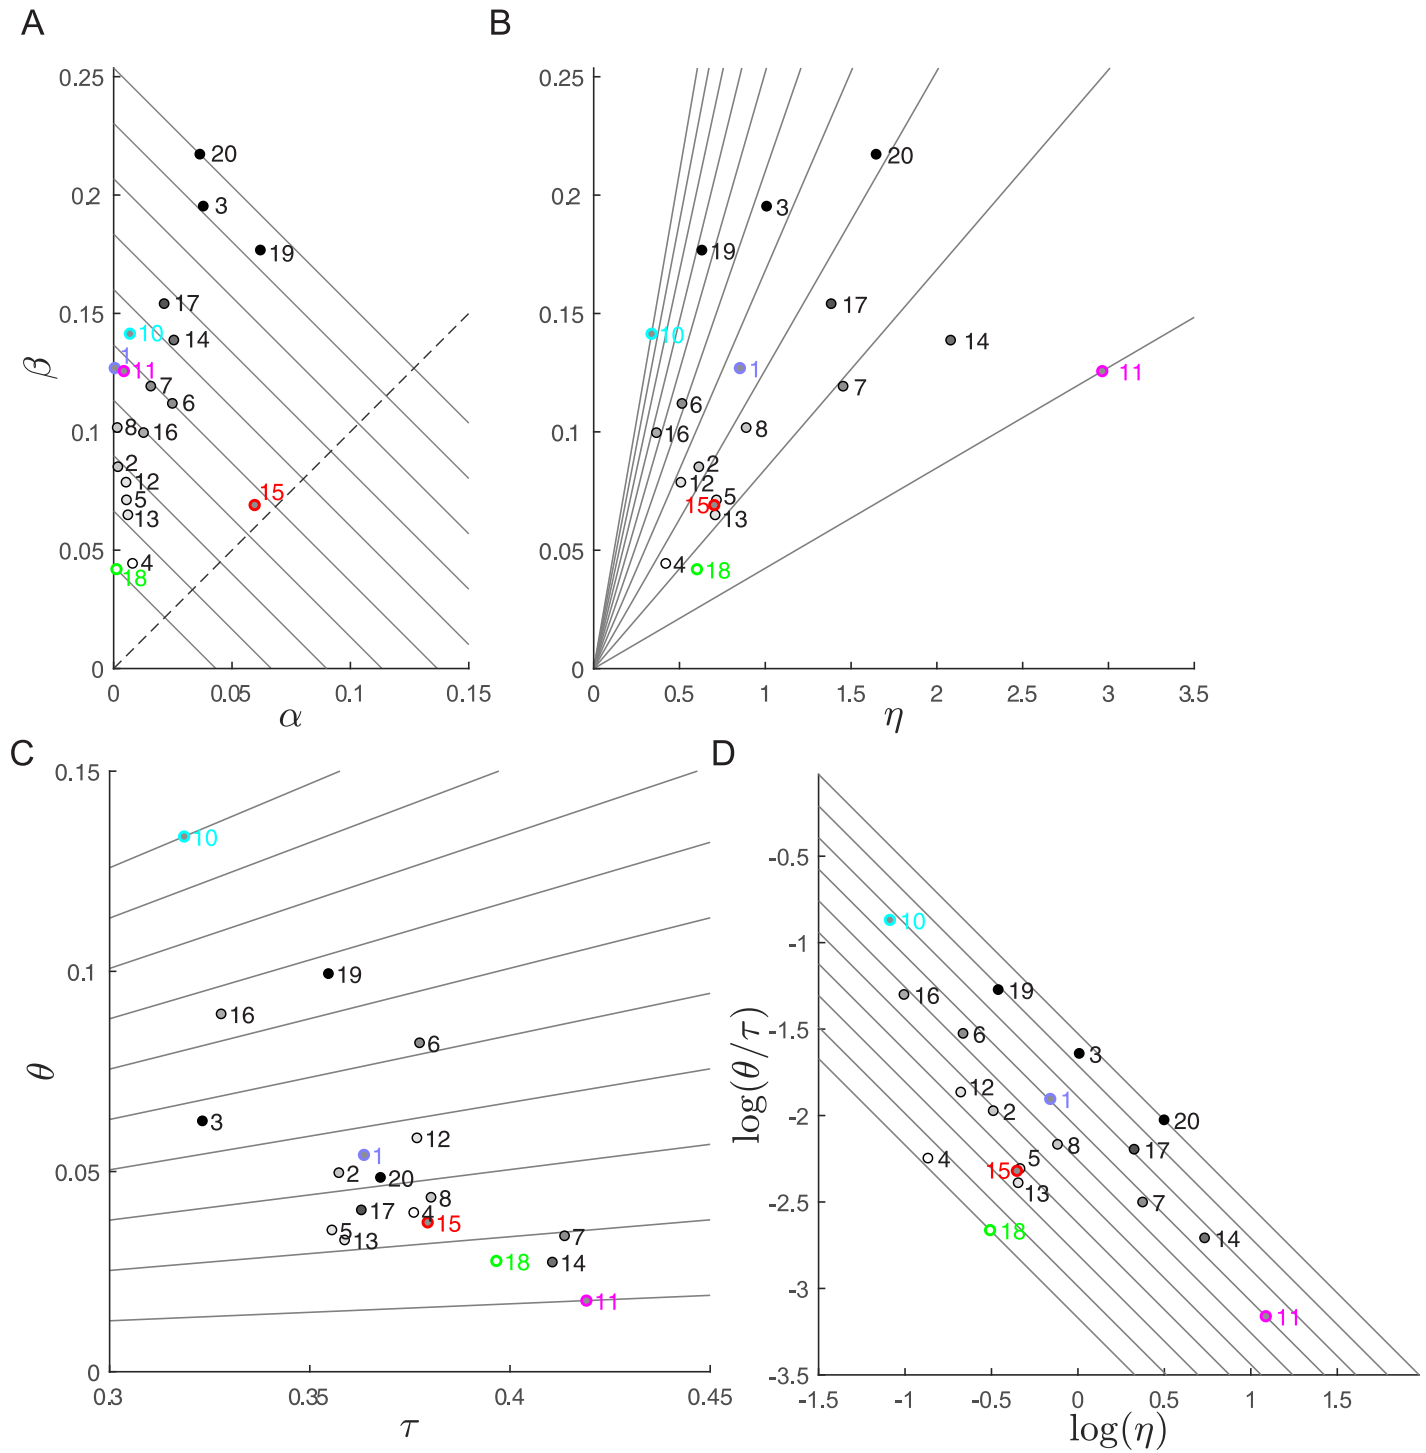


Fig 7. Decomposition of the mean score across participants for target T1. Circular markers are filled with areas that are gray-shaded according to the score (lighter gray for lower mean score). The edge of the markers corresponding to the five participants of Fig 5 are colored with matching colors. To increase the readability of the figure, P9 is not shown. (A) The $\alpha - \beta$ plane and the iso-performance lines. The dashed line represent the direction of maximal change in performance (orthogonal to the iso-performance lines). (B) The $\beta - \eta$ plane. The lines have slopes corresponding to different values of $\frac{\theta}{\tau}$. Participants along the same line, such as P1, P2, and P12, have the similar alignment-to-tolerance ratio and the difference in their β are only due to the noise η . (C) The $\theta - \tau$ (tolerance-alignment) plane. The lines have slopes corresponding to different values of $\frac{\theta}{\tau}$. The negative correlation between θ and τ indicates that participants with mean release action in a region with higher tolerance also tend to have lower alignment with the most noise sensitive directions. (D) The $\log(\eta) - \log(\theta/\tau)$ plane and iso- β lines. As β is the product of θ/τ and η and $\log(\beta) = \log(\theta/\tau) + \log(\eta)$, in the $\log(\eta) - \log(\theta/\tau)$ plane participants with the same β lie on a line with a -1 slope and maximal change of the tolerance-variability index occurs along the orthogonal direction.

<https://doi.org/10.1371/journal.pone.0253626.g007>

color of the edges of the circular markers and show 6D throwing strategies that are similar to those illustrated in the 2D throwing examples of Fig 3, with matching colors.

The $\alpha - \beta$ plane (Fig 7A) characterizes the performance of each participant and the trade-off between bias and tolerance-variability. The *diagonal continuous* lines with slope -1 , corresponding to equal sum $\alpha + \beta$, indicate iso-performing strategies. Participants closer to the origin of the plane are the best performers (e.g., P18 and P4, compare with Fig 1B). Moving away from the origin in the direction of the *dashed* line performance decreases. Participants with the same performance, i.e. on the same diagonal line, differ on their position on the line, thus showing a different trade-off between bias (α) and tolerance-variability (β). For example, P10, P1, and P11 have low α and P15 high α but they are all close to the same iso-performance line and thus have similar performance. In sum, the $\alpha - \beta$ plane shows that our participants differ in their performance because of their specific values of the α and β indicators and that participants with the same performance trade-off α with β .

The $\eta - \beta$ plane (Fig 7B) characterizes the role of the amount of action variability, i.e. the noise η , in determining the tolerance-variability index β . The same contribution to performance due to tolerance-variability may be obtained with very different amounts of variability. For example, P10, P1, and P11 have similar β values but very different η values. This is because larger variability is compensated by lower alignment and higher tolerance. In fact, the slope of the lines in the $\eta - \beta$ plane corresponds to $\frac{\theta}{\tau}$ and participants on lines with lower slope have a lower alignment or higher tolerance.

The interplay between tolerance and alignment is illustrated in the $\tau - \theta$ plane (Fig in Fig 7C). Here the slope of the lines corresponds to a constant $\frac{\beta}{\eta}$, i.e. each line corresponds to a line in the $\eta - \beta$ plane (as $\frac{\beta}{\eta} = \frac{\theta}{\tau}$). While the relationship between τ and θ is not fixed, participants with a higher tolerance also have a lower alignment. This indicates that, when the mean action is located in a more tolerant region of the action space (e.g. higher horizontal velocity, as for P11) it is also possible to further optimize performance by reducing the alignment.

Finally, the $\log(\eta) - \log(\theta/\tau)$ plane (Fig 7D) characterizes the trade-off between the amount of action variability and the other features of the action distribution determining the tolerance-variability index. In the log-log plane the multiplicative contribution of $\frac{\theta}{\tau}$ (determined by the mean action and the alignment of the action covariance) and of the amount of action variability η to β are represented by a sum ($\log(\beta) = \log(\frac{\theta}{\tau}) + \log(\eta)$) and, thus, strategies with equal β are located on lines with slope -1 . Then, this plot clearly shows how participants differs in their tolerance-variability index β (participants with lower β such as P18 and P15 are on lines in the *bottom-left*) and how participants with the same β trade-off noise with tolerance and alignment (e.g. P10 vs. P1 vs. P11).

2.2.3 Consistency of the decomposition of individual action strategies across targets.

To test whether the individual performance-related features of action distributions described by the Hessian-based decomposition are consistent, we compared the α , β , η , τ and θ parameters across both participants and targets. If these parameters are robust indicators of the inter-individual differences in throwing strategy, we expect them to vary significantly across participants but not across targets.

Fig 8 shows the distributions across participants and targets of the parameters of the Hessian-based decomposition in the same combinations of parameters as in Fig 7, which included only target T1. The quadrilaterals (*dashed edge lines* and *color-shaded areas*), with the parameters for the four different targets of each participant as vertexes, are colored according to the mean score rank of each subject. In most cases the parameters do not show large variations across targets, i.e. they represent consistent features that characterize individual strategies. In the $\alpha - \beta$ plane (Fig 8A), a few participants show large across-target variations (e.g. P14, who has

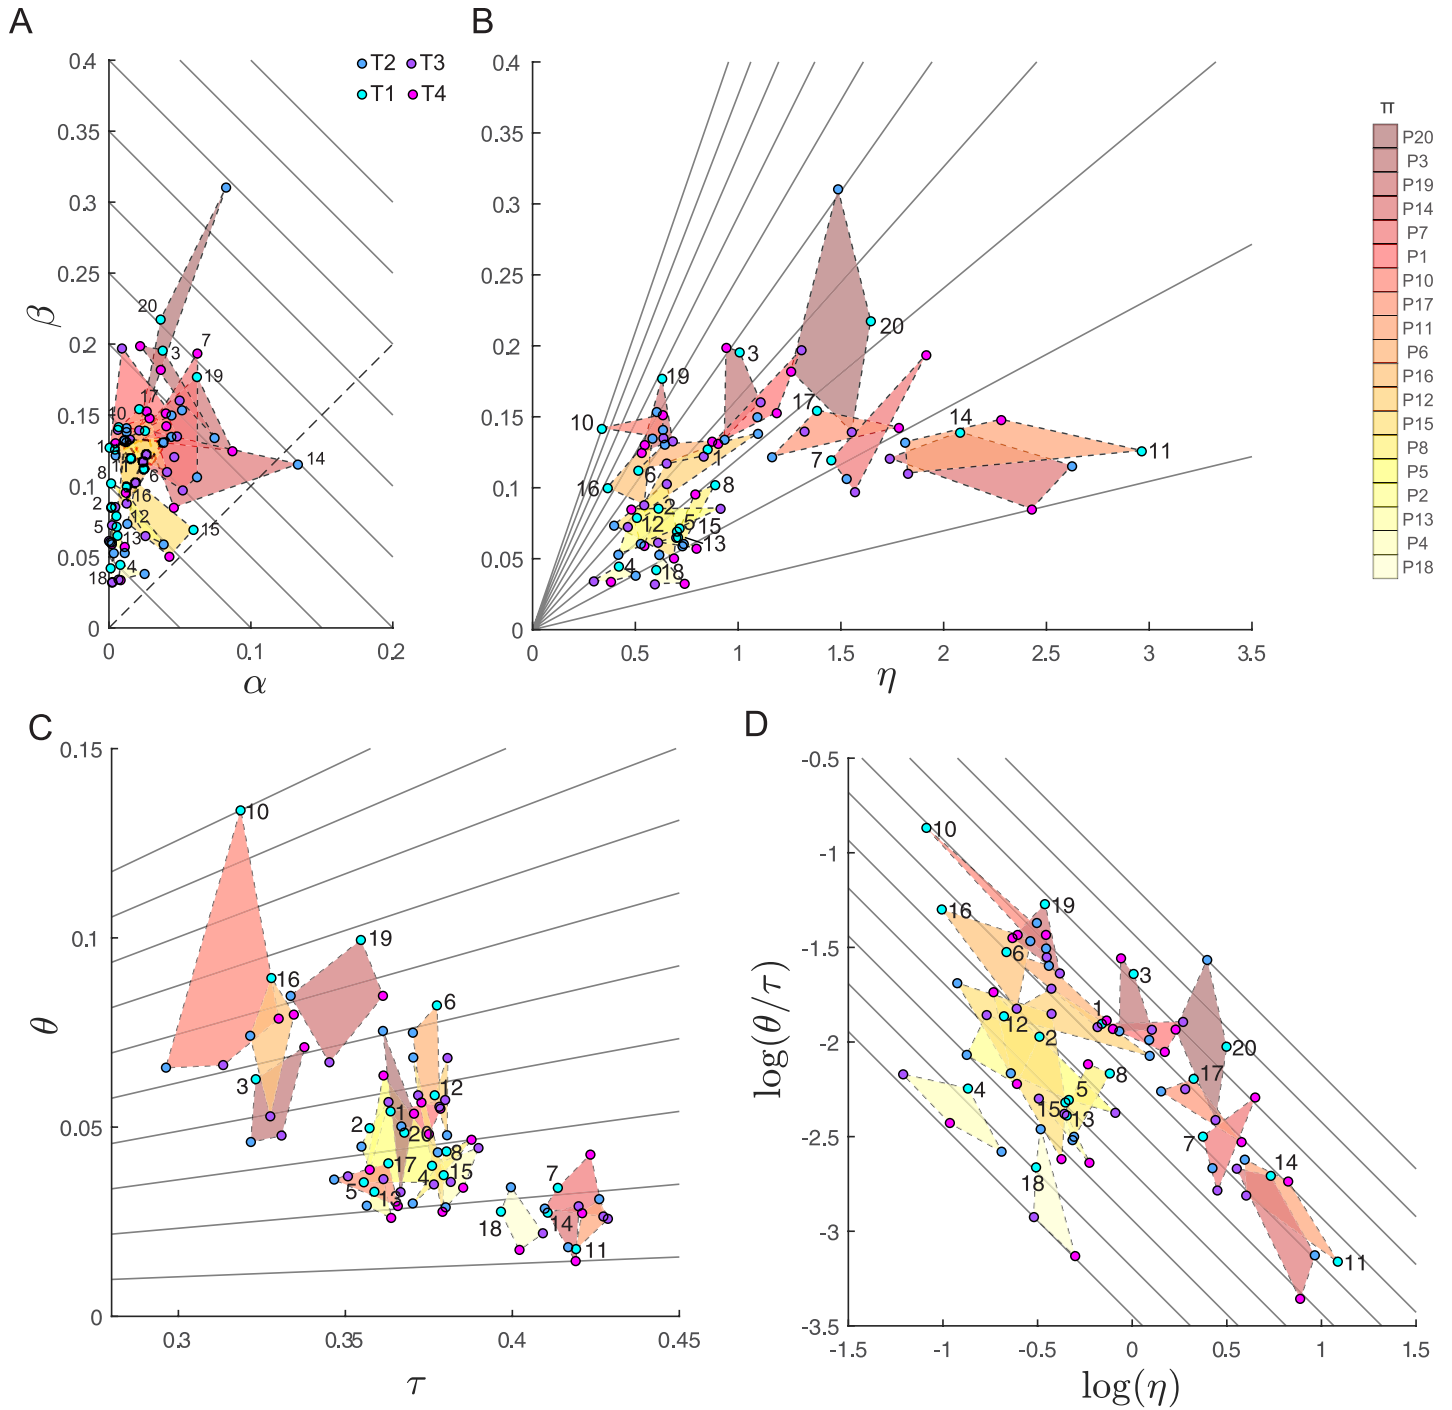


Fig 8. Decomposition of the mean score across participants for all targets. In each panel, the quadrilateral connects the points corresponding to the values of a pair of indicators for the four targets of each participant (excluding P9) and they are colored according to the mean score of that participant (color scale on the top right). The targets are indicated by different colors of the circular markers and the participant number is indicated close to the marker for T1 (except for P7 and P14). (A) The $\alpha - \beta$ plane and the iso-performance lines. (B) The $\beta - \eta$ plane. (C) The $\theta - \tau$ (tolerance-alignment) plane. (D) The $\log(\theta/\tau) - \log(\eta)$ plane.

<https://doi.org/10.1371/journal.pone.0253626.g008>

Table 1. Robustness of decomposition parameters across targets.

parameter	P-value (participant)	P-value (target)
α	***	*
β	***	0.64
θ	***	0.08
τ	***	***
η	***	0.39

P values of a linear mixed-effect model with with participant and target as categorical predictors.

*: $p < 0.05$;

** : $p < 0.01$;

***: $p < 0.001$.

<https://doi.org/10.1371/journal.pone.0253626.t001>

a larger α for T2 than for the other targets; P7, who has a larger β for T4; P20, who has both α and β larger for T2). However, most participants occupy small and distinct regions of the η - β (Fig 8B) and $\log(\eta)$ - $\log(\theta/\tau)$ (Fig 8D) planes. Finally, in the τ - θ plane (Fig 8C) about half of the participants occupy the region with intermediate values of τ .

A linear mixed-effect model with with participant and target identities as categorical predictor (Table 1) of all five parameters supports these qualitative observations, showing that the effect of participant is significant for all parameters while the effect of target is significant only for α ($p = 0.02$) and τ . Thus, α and τ are the only parameters that are not consistent across targets. This is not surprising for τ since it depends on the individual strategy only through the mean release position and velocity, which, in turn, depend on the target position, and, thus, τ varies with the target. Concerning α , the contrast between pairs of targets (fixed effect coefficients) shows that its value is significantly higher for T2 than T1 ($\Delta\alpha = -0.02$, $p < 10^{-2}$) and T3 ($\Delta\alpha = -0.01$, $p = 0.01$), indicating that the bias of a throwing strategies also depends slightly on the target.

3 Discussion

We have developed a novel method to investigate how the distribution of actions in goal directed behaviors relates to individual performance. The method allows to characterize how performance depends on a few critical features of the action distribution, for tasks in which actions are redundant (the same goal may be achieved by multiple actions), high-dimensional (each action is described by a vector with many components) and noisy (actions vary due to stochastic sensory and motor processes). Assuming that the success of an action can be assessed by a scalar cost, i.e. a score, and that such score is a smooth function of the action, we derived an approximate analytical relationship between the mean score and the first two moments of the actions distribution: the mean action $\bar{\mathbf{a}}$ and action covariance Σ^a across multiple trials. Performance, defined as the mean score, can be approximated as the sum of two components: the score of the mean action ($\alpha(\bar{\mathbf{a}})$) and a tolerance-variability index ($\beta(\bar{\mathbf{a}}, \Sigma^a)$). The α parameter, when different from zero, measures deviations of the mean action from the set of actions that accurately achieve the goal (solution manifold). The β index, instead, measures how the mean score is affected by actions variability and by the geometry of the action-to-score function (determining the sensitivity to noise as a result of the non-linearities of the action-to-score function around the mean action and their alignment with the directions of largest variability in the action distribution). Such index results from the product of three terms: (i) the *noise* (η), computed as the sum of the variances of the individual components of the action vector (i.e. the trace of the action covariance matrix); (ii) the *tolerance* of the score

function (τ), quantifying the overall sensitivity of the score to deviations from the mean action due to the curvature of the action-to-score function captured by the Hessian; (iii) the *alignment* (θ), a scalar measure of the effect on mean score of the relative orientation between the most sensitive directions of the action-to-score function and the directions of maximum variability of the actions. Thus, these five indicators provide a compact yet informative characterization of the features of the action distributions that affect performance in relation to a specific task, and allow to capture detailed facets of individual strategies in goal directed behaviors.

We have applied this method to characterize individual performance and variability in unconstrained overarm throwing actions of twenty non-trained participants. Across participants there were remarkable inter-individual differences in the α , β , η , τ and θ parameters (Fig 7) and most of these differences were consistent across targets (Fig 8). In line with previous works focusing on low-dimensional throwing tasks [7, 17], in our unconstrained high-dimensional throwing task we found that skilled participants have small α (accurate mean action) and small β (tolerance-variability index). Still, it is possible to further differentiate different optimizing strategies, as low β can be achieved by minimizing action variability or by compensating the higher variability in action execution (η) with higher tolerance τ and smaller alignment (i.e., smaller θ). Moreover, examining the scatter plot of different pairs of parameters of individual participants, we identified specific combinations of parameters that do not affect performance, but correspond to different throwing strategies. First, similar performance can be achieved trading off bias (α) with tolerance-variability (β). Second, similar β can be achieved trading off alignment and tolerance (θ/τ) with variability (η). Such interplay between variability and geometric features of the score can be observed in several pairs of action variables (e.g. in the v_y - v_z plane, Fig 5) but are characterized, fully yet compactly, by the indicators of the decomposition (Fig 7).

3.1 Comparison with related approaches

In redundant, high-dimensional, and noisy tasks it is not enough to characterize the mean and the covariance of the action distribution to fully capture the relationship between action variability and performance. In agreement with earlier computational approaches addressing variability in multivariate actions [7, 15–17], our method highlights the key role of the geometry of the action-to-score function (captured by the Hessian) to assess how action variability affects performance. Differently from first-order methods such as UCM, GEM, or the more recent approach in [24], which characterize the local geometry with a linear approximation (expressed through the Jacobian matrix or the gradient vector), our method relies on a second-order approximation (based on the Hessian matrix). The main reason for which our method does not depend on the first-order term of the Taylor expansion of the action-to-score function and requires a second-order approximation is the fact that we are considering the mean score rather than the variability in action or outcome space as a measure of performance. As indicated in Eqs (14) and (15), the mean score does not depend on the gradient of the action-to-score function computed at the mean action, the reason being that the first-order term of the expansion is multiplied by the mean deviation from the mean action, which is null by definition. In other terms, changes in score (with respect to the mean) associated to actions that deviate from the mean action sum up to zero in the linear approximation of the action-to-score function. Indeed, for a linear action-to-score mapping the mean score is given simply by the score of the mean action, as all higher order derivatives in the expansion are null. Thus, in a quadratic approximation, it is only the local curvature of the action-to-score function, captured by the Hessian matrix, that affects the mean score.

How action distribution affects the mean score in a goal-directed behavior has been addressed by the TNC methods [7, 11, 13, 20]. The methods have been developed for, and applied to, a two-dimensional throwing task inspired by the skittle game, in which participants have to hit a target by releasing through a rotating joint (i.e., the action parameters are the release angle and tangential velocity) a virtual ball that could rotate in a plane around a pole. The TNC methods take into account the geometry of the action-to-score mapping implicitly, by evaluating the effects on performance of different action distributions through surrogate data (when comparing pairs of distributions [7]) or by exhaustive search (when comparing with optimal distributions [20]). In contrast, our method explicitly decomposes the contribution of different features of the action distribution through a Taylor expansion. Such analytical approach overcomes the disadvantage of the TNC methods concerning the use of numerical procedures necessary to generate surrogate data or to search for optimal distributions, which limits its applicability to high-dimensional actions. Moreover, our method allows to determine the contribution of the local geometry and of the action variability independently on each other. As illustrated in the 2D throwing examples of Fig 4A and detailed in the Appendix (see Section A in S1 Appendix), under the assumption of a smooth action-to-score function, for which the Hessian matrix is well defined, the tolerance, noise, and covariation terms of the TNC decomposition correspond to specific combinations of the terms in our decomposition. However, importantly, all the three terms of the TNC decomposition depend on both the Hessian and the action covariance, thus they do not independently characterize the contributions of the local action-to-score geometry, of the total action variability, and of the alignment between action variability and score tolerance τ , η , and θ . Moreover, when assessing a single distribution, the TNC-Cost analysis does not provide a unique decomposition of performance, as the three costs do not sum up to the mean score. Then, it is not possible to directly characterize the trade-off between the different components of the variability and local task geometry as with the Hessian-based decomposition. As an advantage, however, the TNC method does not rely on any assumption on the action-to-score function, such as smoothness and adequateness of a second-order approximation.

A fundamental difference between GEM and both the TNC and our approach, is the choice of performance measure. In the study of goal-directed motor skills performance is typically associated with a measure of *error* with respect to a predefined goal. Reduction of task errors and their variability is broadly recognized as an indicator of skilled performance. However, more recent views of human motor control, based on decision-making theory, propose an alternative definition of performance which is based on the concept of a score/loss function that essentially assigns a number (or a cost) to a given task error. In this perspective, the *mean* or *expected loss* is taken as measure of performance over repeated trials of a given motor task. Central to the GEM analysis is a goal function $e(\mathbf{a}) = \mathbf{x}_T - f(\mathbf{a})$ that expresses the error between a desired goal/target \mathbf{x}_T and the outcome $f(\mathbf{a})$ of a given action \mathbf{a} . By linearizing this error/goal function around the mean action of a strategy, GEM quantifies the overall contribution of *tolerance*, *noise* and *alignment* with the task/goal-level variability, i.e. $var(\mathbf{e})$. The Hessian-based decomposition proposed in this work extends the GEM analysis to decomposition of the mean of a performance indicator. In particular, we have shown that is the Hessian and not the Jacobian that affects the expected score/loss around the mean action of a strategy (see Eq (30)).

3.2 Assumptions and limitations

Our decomposition method requires a smooth action-to-score function and it provides an accurate estimate of the mean score only if the non-linearities in such function are adequately approximated by the second-order term of the Taylor expansion over the domain spanned by

the actions. The assumption of smoothness (or at least continuity of the function and all partial derivatives up to the second order) is valid for a broad class of score functions, such as most penalty or reward functions usually employed to quantify task performance. For goal-directed tasks, requiring to minimize the distance from a target, the squared distance is a good choice because it leads to an action-to-score function which is twice differentiable everywhere the action-to-outcome function is smooth. The squared distance is preferable over the Euclidean distance because the latter has a singularity in the second derivative at zero, i.e. on the solution manifold. However, if the subject is attempting to minimize (maximize) the score, as the distance and the squared distance have the same minimum (maximum), both functions capture the control strategy equally well.

Another key assumption in our approach is that the second-order Taylor expansion of the action-to-score function around the mean action provides an acceptable approximation. As shown in Fig 6, for almost all participants and targets, the estimation of the mean score based on such approximation ($\alpha + \beta$) is close to the actual mean score ($E[\pi]$). The only exception is participant P9 who had a poor performance and a very large variability in the ball release parameters. Indeed, the validity of the quadratic approximation depends on the nature of the non-linearities of the action-to-score function and the range of the deviations from the mean, i.e. from the relative spatial scales characterizing the concentration of the action distribution and the Hessian. Thus, if behavior is very erratic, our decomposition may become inaccurate for the entire set of actions and may be restricted to a more concentrated subset. However, considering that participants in our sample were untrained throwers, it is noticeable that the quadratic approximation was good for all but one out of twenty participants. This suggests that our methods could be safely applied to more controlled tasks, e.g. in evaluating athletes performances (as athletes do not typically exhibit high variability in motor actions) or in assessing motor skill learning, where training tends to quickly reduce motor variability. However, in cases where the variability is high, such as during the initial exploration phase when learning a novel motor skill, one could apply the decomposition on local clusters of data and still provide a compact characterization through the parameters of Hessian-based decomposition of each cluster. We plan to develop such approach based on a mixture of Hessian-based decompositions and to apply it to the investigation of throwing skills learning in future work.

Our decomposition method relies on the computation of the action covariance and the Hessian of the action-to-score function. These matrices and some of the parameters of the decomposition depend on the choice of the coordinate system in action space. In particular, the noise η and the tolerance τ , being defined as traces of the covariance and Hessian matrices, respectively, change under coordinate transformations (unless a metric is chosen [25]). However, the α term is a scalar (i.e. is a single number corresponding to the score associated to the mean action) and it does not depend on coordinates. The β term is the trace of the product of the covariance and the Hessian matrices and it is invariant under affine coordinate transformations, given that Σ^a and H^a transform in opposite ways (see Section B in S1 Appendix). Thus, re-scaling of positional and velocity coordinates due to different choices of measurement units does not affect the decomposition of mean score as a sum of α and β . However, β is not invariant, in general, for non-linear coordinate transformations, such as the transformation from Cartesian to polar coordinates. Indeed, the dependence on action coordinates has raised concerns about the reliability of the TNC decomposition [26, 27] and of UCM and GEM methods [28]. While such dependence may provide an opportunity to evaluate the role of different coordinate systems for control [29], it has also been noticed that geometric properties of the action-to-score function such as the solution manifold do not depend on coordinates [28]. In our decomposition, if the mean of the action distribution is on the solution manifold ($\alpha = 0$), β is invariant also under non-linear transformations, because the non-linear term in the

transformation of H^a depends on the gradient of the action-to-score function, which is null on the solution manifold. Moreover, if the action distribution is not centered on the solution manifold but it is concentrated (i.e. η is small) the change in β due to non-linear coordinate transformations may be negligible.

A further limitation of our approach, which is shared with GEM and TNC, is that the decomposition cannot reveal the temporal structure of inter-trial fluctuations, as multiple trials are needed to compute the mean and the variance of an individual's action strategy. Recent work addresses this issue with an inter-trial error correction model that predicts both the temporal and geometric structure of variability near the goal equivalent manifold (GEM) of a simplified shuffle board task [30]. Furthermore, the variability analysis is shown to be coordinate-independent as the characterization is performed in the eigenspace of the error-correcting controller matrix.

3.3 Applications to motor skill learning

In this work we have focused on characterizing steady-state performance and individual action distribution during short experimental sessions rather than on skill improvement over multiple sessions. Future work will include longitudinal studies to understand if and how the observed inter-individual differences are related to the time course and the magnitude of individual performance improvements and skill learning. Current theories of human sensorimotor control suggest the existence of two distinct mechanisms underlying motor skill learning: a model-based system that improves motor performance guided by an internal forward model of the body and the environment, which is updated based on prediction errors [31]; and a model-free system in which learning is driven by reinforcement and punishment of successful/erroneous actions [32, 33]. Motor adaptation studies, in which a systematic perturbation of the environment is introduced by means of force fields of visuomotor rotations, suggest that the model-based system is responsible for the quick adaptation/compensation of the mean error. The model-free system, driven by reinforcement and punishment, regulates instead motor variability, and is hence responsible for the slow reduction of the variable errors. However, the interplay between these two learning mechanisms, remains poorly understood.

In analyzing our free overarm throwing data, we have highlighted the existence of iso-performing participants, such as *P1*, *P11* and *P15*, which have the same mean score, but different contributions of α and β . Do inter-individual differences in terms of α and β translate into individual differences in terms of performance improvement? In future work we plan to use the proposed framework to study the acquisition of throwing skills in virtual reality environments in which we can alter both the dynamics of the ball, for instance by manipulating the (virtual) gravity field, as well as the task score geometry, in this work assumed quadratic and isotropic in both task directions. As adapting to an altered dynamics requires learning a new forward model while a new task geometry changes the reward function, the dissociation between these two contributions might allow us to dissociate between model-based and model-free learning and to understand how initial inter-individual differences in terms of performance, variability and score tolerance translate into individual performance improvement.

3.4 Summary and conclusions

We have introduced a novel method to characterize the key features of the distribution of high-dimensional and redundant actions that affect performance, defined as the mean of the score assigned at individual actions. We have applied the method to the investigation of inter-individual differences in unconstrained throwing. We found that the indicators derived from the Hessian-based decomposition allow to identify specific and consistent features relating

individual throwing strategies to different performance level and to understand how different strategies achieve similar performances. Participants differed in their throwing performance because they consistently differed in either the score of their mean action (bias, α) or the level of their tolerance-variability index (β), a measure of the interplay between action variability and action-to-score geometry. The same performance could be achieved trading off α with β and the same β could be achieved trading off θ/τ with η . In sum, the compact characterization of the relation between high-dimensional, redundant, and noisy action distributions and performance provided by our Hessian-based decomposition may be applied to a variety of complex real-life motor skill, opening up new opportunities both for systematic investigations of inter-individual differences in real-life motor skills and for practical applications to training of complex motor tasks. As motor control investigators, we plan to address in future studies how the individual parameters of the Hessian-based decomposition relate to individual motor learning capabilities, a first step to unravel the fundamental mechanisms underlying individual learning differences. A sport trainer could also use the Hessian-based decomposition indicators of a trainee, based on the analysis of a few tenths of repetitions of a specific motor skill, to select an optimal individualized training strategy according to those indicators.

4 Methods

4.1 Performance as expected value of non-linear and high-dimensional action scores

Let us assume that, at every trial t , an individual generates an action $\mathbf{a}_t \in A \subset R^n$ with some random noise such that the action distribution can be described by a probability density function (p.d.f.) p_A . Let us also assume that, at every trial, the action receives a score point $\pi_t = s^a(\mathbf{a}_t)$ through a *score* (or loss) function $s^a: \mathbf{a} \in R^n \rightarrow R^+$ which punish motor actions according to some optimality criteria such as task errors or metabolic cost of an action. For instance, in Fig 1, the score function assigns a penalty score to an action, that is the squared distance between the action outcome $\mathbf{x}(\mathbf{a})$ and the target position \mathbf{x}_T . We define the *performance* of a motor strategy p_A as the *expected score*:

$$E[\pi] = \int_{A \subset R^n} s^a(a_1, a_2, \dots, a_n) p_A(a_1, a_2, \dots, a_n) da_1 da_2 \dots da_n. \quad (10)$$

which is the sum over all possible actions of the probability of a taking a given action times its score (integrated in R^n). In this framework, an individual motor strategy p_A is optimal if, over repeated attempts, minimizes the *average* score. Similarly, given two motor strategies we can establish if they are *equivalent* or if one is better than the other by simply comparing their expected score. In practice, the score function may be highly non-linear and the action space high dimensional ($n \gg 1$) making difficult to find an analytic solution to (10). Therefore, we do not usually solve (10) but instead approximate the expected score with its sample average as in Fig 1B, where best performers, such as P18 and P4 have the lowest average quadratic error. Nevertheless, the sample average by itself is not informative enough to understand *why* the individual action distributions of P18 and P4 are much more performing compared to the majority of participants.

4.1.1 Quadratic approximation of the expected action score. By restricting the action score to the class of continuous and at least twice differentiable functions, and assuming that the action distribution is sufficiently localized around the mean action, it is possible to find an approximate but analytic solution to (10).

Let's assume that an individual selects motor actions according to a p.d.f. with *expected or mean action* $E[\mathbf{a}] = \bar{\mathbf{a}} \in R^n$, and covariance matrix $\Sigma^a \in P^n$ (symmetric and positive definite).

In other words, at every trial t , an individual generates an action \mathbf{a}_t according to the following stochastic model:

$$\mathbf{a}_t = \bar{\mathbf{a}} + \boldsymbol{\delta}_a \tag{11}$$

where $\boldsymbol{\delta}_a \in R^n$ is the *stochastic component* of the individual strategy, which, at every trial, ‘perturbs’ the mean action $\bar{\mathbf{a}}$. The mean action represents an individual preference in choosing, *on average*, a given action and the covariance matrix $\Sigma^a = E[\boldsymbol{\delta}_a^T \boldsymbol{\delta}_a]$ represents action covariation/correlation (*action variability*) across multiple trials.

Let us also assume, that locally, i.e. in a neighborhood of the average action $\bar{\mathbf{a}}$, the score $s^a(\mathbf{a}_t)$ of an action \mathbf{a}_t , can be approximated with the following second-order Taylor expansions:

$$s^a(\bar{\mathbf{a}} + \boldsymbol{\delta}_a) \approx s^a(\bar{\mathbf{a}}) + [\nabla_a s^a(\bar{\mathbf{a}})]^T \boldsymbol{\delta}_a + \frac{1}{2} \boldsymbol{\delta}_a^T H(\bar{\mathbf{a}}) \boldsymbol{\delta}_a \tag{12}$$

where $\nabla_a s^a(\bar{\mathbf{a}}) = \left[\frac{\partial s^a}{\partial a_1}, \frac{\partial s^a}{\partial a_2}, \dots, \frac{\partial s^a}{\partial a_n} \right]$ is the gradient of the action score function evaluated at the average action $\bar{\mathbf{a}}$ and:

$$H_{ij}(\bar{\mathbf{a}}) = \frac{\partial^2 s^a}{\partial a_i \partial a_j}(\bar{\mathbf{a}}). \tag{13}$$

is the $n \times n$ Hessian matrix (assumed to be symmetric and positive definite) of the action score evaluated at $\bar{\mathbf{a}}$.

Inserting (12) in the integral (10), we can write the expected action score as the sum of three terms:

$$E[\pi] \approx E[s^a(\bar{\mathbf{a}})] + [\nabla_a s^a(\bar{\mathbf{a}})]^T E[\boldsymbol{\delta}_a] + E\left[\frac{1}{2} \boldsymbol{\delta}_a^T H(\bar{\mathbf{a}}) \boldsymbol{\delta}_a\right] \tag{14}$$

the first term $E[s^a(\bar{\mathbf{a}})]$ is simply *the score of the average action* $s^a(\bar{\mathbf{a}})$ given that the expected value of a constant is the constant itself. The second term, $[\nabla_a s^a(\bar{\mathbf{a}})]^T E[\boldsymbol{\delta}_a]$, vanishes whenever the quadratic approximation is evaluated at the mean action $\bar{\mathbf{a}}$, given that in such condition $E[\boldsymbol{\delta}_a] = 0$. The last term corresponds to the expected value of a *quadratic form*, which is well known to be equal to $trace\left(\frac{1}{2} H(\bar{\mathbf{a}}) \Sigma^a\right)$ [34]. This term is zero: i) when the score is a linear function of the action, in which case the Hessian is zero, ii) when actions are not stochastic $\Sigma^a = 0$, or when H and Σ^a are ‘orthogonal’. In all other cases this term will influence the expected action score.

Under such hypothesis the expected action score (10) can be (locally) approximated as:

$$E[\pi] \approx s^a(\bar{\mathbf{a}}) + trace\left(\frac{1}{2} H(\bar{\mathbf{a}}) \Sigma^a\right) = \alpha(\bar{\mathbf{a}}) + \beta(\bar{\mathbf{a}}, \Sigma^a) \tag{15}$$

where:

$$\alpha(\bar{\mathbf{a}}) \equiv s^a(\bar{\mathbf{a}}) \tag{16}$$

is *the score of the average action*, and

$$\beta(\bar{\mathbf{a}}, \Sigma^a) \equiv trace\left(\frac{1}{2} H(\bar{\mathbf{a}}) \Sigma^a\right) \tag{17}$$

is the *tolerance-variability index* which captures how local non-linearities of the action score ($H(\bar{\mathbf{a}})$) and variability in the action strategy (Σ^a) affect (increase/decrease) performance, i.e.

the mean score $E[\pi]$. Note that the definition of β corresponds to half the Frobenius inner product between the Hessian and the Covariance matrix. In the next section we show that it can be decomposed into three independent components: *the noise, the score tolerance* and the *alignment*.

4.2 Hessian-based decomposition of motor skills

4.2.1 Principal variability directions and noise η . The action covariance matrix Σ^a , symmetric and positive-definite, can be decomposed (principal component analysis) into singular values (or eigenvalues) and singular vectors (or eigenvectors):

$$\Sigma^a = U^\Sigma \Lambda^\Sigma U^{\Sigma T} = [\mathbf{u}_1^\Sigma, \mathbf{u}_2^\Sigma, \dots, \mathbf{u}_n^\Sigma] \text{diag}(\lambda_1^\Sigma, \lambda_2^\Sigma, \dots, \lambda_n^\Sigma) [\mathbf{u}_1^\Sigma, \mathbf{u}_2^\Sigma, \dots, \mathbf{u}_n^\Sigma]^T. \tag{18}$$

with $\lambda_1^\Sigma > \lambda_2^\Sigma > \dots > \lambda_n^\Sigma$. The larger a singular value λ_j^Σ , the more variable will be the action strategy along its associated *principal variability direction* \mathbf{u}_j^Σ .

We define the *noise* η of an individual strategy as the total variation of the action distribution:

$$\eta(\Sigma^a) \equiv \text{trace}(\Sigma^a) = \sum_{j=1}^n \sigma_{jj} = \sum_{l=1}^n \lambda_l^\Sigma \tag{19}$$

4.2.2 Principal sensitivity directions and tolerance τ . When locally, i.e. around the mean action $\bar{\mathbf{a}}$, the score is a continuous and at least twice differentiable function of the action variables, the Hessian is an $n \times n$ symmetric matrix (Schwarz’s theorem) which can be written as:

$$H = U^H \Lambda^H U^{HT} = [\mathbf{u}_1^h, \mathbf{u}_2^h, \dots, \mathbf{u}_n^h] \text{diag}(\lambda_1^H, \lambda_2^H, \dots, \lambda_n^H) [\mathbf{u}_1^h, \mathbf{u}_2^h, \dots, \mathbf{u}_n^h]^T. \tag{20}$$

The diagonal matrix $\Lambda^H = \text{diag}(\lambda_1^H, \lambda_2^H, \dots, \lambda_n^H)$ contains the n singular values $\lambda_1^h, \lambda_2^h, \dots, \lambda_n^h$ (with $\lambda_1^h \geq \lambda_2^h \geq \dots, \lambda_n^h$) of the Hessian matrix. The orthonormal matrix U^H contains the associated singular vectors $[\mathbf{u}_1^h, \mathbf{u}_2^h, \dots, \mathbf{u}_n^h]$. The larger a singular value λ_j^H , the greater will be the change of the score (i.e. the greater the curvature will be) along its associated *principal sensitivity direction* \mathbf{u}_j^h . When the Hessian matrix is positive-definite, i.e. the singular values λ_i^H are positive, the score function is locally convex and therefore the mean action is in, or ‘close to’, a (local/global) minimum of the score function. Conversely, negative eigenvalues are representative of concave regions of the score function, while eigenvalues with mixed signs suggest that the average action is in/close-to a saddle point of the score function. In this work we will focus on score functions which are locally convex and for which the Hessian matrix is semi positive-definite, i.e. all eigenvalues are greater or equal than zero, although what follows can be generalized to more complex score functions having a landscape with many minima, maxima and saddle points.

The eigenvalues λ_i^H express the *sensitivity* of the score to stochastic perturbations. The larger λ_i^H , the more sensitive the score function is to perturbations $\delta_{\mathbf{a}}$ which are directed along the i -th principal sensitivity vector \mathbf{u}_i^h . As a local, scalar measure, of ‘total curvature’ of the score, we define the *sensitivity* of the score as the $\text{trace}(H) = \sum_{i=1}^n h_{ii} = \sum_{l=1}^n \lambda_l^H$.

The local *score tolerance* then, is defined as the reciprocal of the score sensitivity:

$$\tau(\bar{\mathbf{a}}) \equiv \frac{1}{\text{trace}(H(\bar{\mathbf{a}}))}. \tag{21}$$

4.2.3 Alignment θ . With the above definitions of tolerance and noise, by normalizing both the Hessian and the covariance matrix by their respective traces, we can rewrite the tolerance-variability index β as:

$$\beta = \text{trace}\left(\frac{1}{2} \frac{H}{\text{trace}(H)} \frac{\Sigma}{\text{trace}(\Sigma)}\right) \text{trace}(H) \text{trace}(\Sigma) = \text{trace}\left(\frac{1}{2} \bar{H} \bar{\Sigma}\right) \frac{\eta}{\tau} = \frac{\theta}{\tau} \eta \tag{22}$$

where the *alignment* θ :

$$\theta(\bar{\mathbf{a}}, \Sigma^a) = \text{trace}\left(\frac{1}{2} \bar{H} \bar{\Sigma}\right) \tag{23}$$

is a scalar that measures the relative orientation between (normalized) principal sensitivity vectors and (normalized) principal variability vectors. In other words, the more the directions of maximal variability \mathbf{u}_i^Σ are aligned with the directions of maximal sensitivity \mathbf{u}_i^H of the score, the larger the effect of variability on the β and hence on the expected action score.

In conclusion, the performance of a motor skill, or its expected action score can be approximated (and decomposed) as:

$$E[\pi] \approx \alpha(\bar{\mathbf{a}}) + \beta(\bar{\mathbf{a}}, \Sigma^a) = \alpha(\bar{\mathbf{a}}) + \frac{\theta(\bar{\mathbf{a}}, \Sigma^a)}{\tau(\bar{\mathbf{a}})} \eta(\Sigma^a) \tag{24}$$

4.3 Application to throwing tasks

Throwing skills, as many other motor skills, are usually assessed by means of score functions which essentially define the objective of the throwing task. For instance, for a javelin thrower, the score may be a function of the longitudinal distance travelled by the javelin. The further the javelin lands, the larger the score assigned to the throwing action. Conversely, for a dart thrower, the goal is not to throw the dart as far as possible, but as accurate as possible, and hence, as in Fig 1A, the score function could assign a penalty increasing with the distance between the landing position of the projectile and the center of the target.

It should be noted that in our experimental protocol [22] participants did not receive any explicit performance feedback (or score) at end of each throwing trial (but they could see the arrival position of the ball on the target board) and therefore, in this work, in line with computational and experimental evidences [1, 31], we assume that participants optimize an accuracy score which penalizes the squared error between the outcome \mathbf{x} of a release action and the target position \mathbf{x}_T [31]:

$$\pi = s^x(\mathbf{x}, \mathbf{x}_T) = (\mathbf{x} - \mathbf{x}_T)^T (\mathbf{x} - \mathbf{x}_T) \equiv e^2. \tag{25}$$

Written in this form, the accuracy score, represents a *task score* $s^x: R^2 \rightarrow R^+$ which penalizes the two-dimensional action outcomes \mathbf{x} with a scalar score π . To find the relationship between release actions and quadratic (task) error, i.e. the *action-to-score function* $s^a: \mathbf{a} \in R^6 \rightarrow R^+$, we need to express the task score (25) as a function of the release parameters.

Assuming a point-mass projectile and hence neglecting friction and Magnus’s forces, the projectile trajectory $f(\mathbf{a}, t)$ can be predicted from the release parameters $\mathbf{a} = [\mathbf{p}_0, \mathbf{v}_0]$ as:

$$f(\mathbf{p}_0, \mathbf{v}_0, t) = \begin{cases} x(t) = x_0 + v_{x0}t \\ y(t) = y_0 + v_{y0}t \\ z(t) = z_0 + v_{z0}t - \frac{1}{2}gt^2 \end{cases} \tag{26}$$

with $\mathbf{p}_0 = [x_0, y_0, z_0]^T$ and $\mathbf{v}_0 = [v_{x0}, v_{y0}, v_{z0}]^T$.

For a target board oriented as in Fig 1A, i.e. with the normal pointing in the longitudinal direction y , the time of impact of the ball with board can be estimated as $t_I = \frac{y_b - y_0}{v_{y0}}$, i.e. as the ratio between the longitudinal distance of the projectile with the board at release (y_b is the coordinate of the board with respect to the world frame) and the velocity of the ball along such direction (that is constant according to the model in (26)). Hence, at the time of impact, the projectile will hit the board at:

$$\mathbf{x}_i = f(\mathbf{a}, t_I) = \begin{cases} x_0 + v_{x0} \frac{y_b - y_0}{v_{y0}} \\ y_b \\ z_0 + v_{z0} \frac{y_b - y_0}{v_{y0}} - \frac{1}{2}g \left(\frac{y_b - y_0}{v_{y0}} \right)^2 \end{cases} \tag{27}$$

Substituting the above system of equations into (25), let us writing the *action score* as:

$$e^2 = s^x(\mathbf{x}_i, \mathbf{x}_T) = s^x(f(\mathbf{a}, t_I), \mathbf{x}_T) = s^a(\mathbf{a}, \mathbf{x}_T) \tag{28}$$

i.e. as a scalar function $\mathbf{a} \in R^6 \rightarrow R^+$ of the throwing action \mathbf{a} :

$$e^2 = s^a(\mathbf{a}, \mathbf{x}_T) = \left(x_0 + v_{x0} \frac{y_b - y_0}{v_{y0}} - x_T \right)^2 + \left(z_0 + v_{z0} \frac{y_b - y_0}{v_{y0}} - \frac{1}{2}g \left(\frac{y_b - y_0}{v_{y0}} \right)^2 - z_T \right)^2 \tag{29}$$

Given an individual release strategy with mean action $\bar{\mathbf{a}}_0$ and covariance Σ^a and aiming at hitting a desired target \mathbf{x}_T , the mean squared error ($E[e^2] = E[\pi]$) can be approximated with (15):

$$E[\pi] = E[e^2] \approx \alpha(\bar{\mathbf{a}}) + \beta(\bar{\mathbf{a}}, \Sigma^a) = \alpha(\bar{\mathbf{a}}) + \frac{\theta(\bar{\mathbf{a}}, \Sigma^a)}{\tau(\bar{\mathbf{a}})} \eta(\Sigma^a) \tag{30}$$

where $\alpha(\bar{\mathbf{a}})$ is just (29) evaluated at $\bar{\mathbf{a}}$, i.e. the quadratic error of the outcome of the mean action, and β can be decomposed into the three components η , τ , and θ by using the covariance matrix of the action strategy Σ^a and the 6×6 Hessian of (29) evaluated at $\bar{\mathbf{a}}$.

In this work, the 6×6 Hessian matrix of (29) is calculated with the Matlab Symbolic Toolbox for each target condition and for each individual strategy.

4.4 2D throwing simulation

To illustrate the Hessian-based decomposition in simple task with 2D actions, we considered a toy model of throwing. We assumed that a projectile is released from a fixed position and moves in a vertical plane to hit a target on a vertical axis at a given distance from the release position (Fig 3A). Thus, the action \mathbf{a} corresponds to a 2D vector whose components are the

horizontal release velocity (v_{y_0}) and the vertical release velocity (v_{z_0}), the outcome x corresponds to the arrival height on the vertical axis, and the score the squared distance from the target (x_T).

We simulated distributions of release actions with a given mean \bar{a} and covariance Σ_a (Figs 3 and 4) by drawing 50 random samples from a multivariate Gaussian distribution. We simulated a target at a distance of 5 m from the release position and at height 0.7 m below it, reproducing the mean conditions for target T1 in our sample of 6D throwing participants (see below). In addition to decomposing the mean score according to Eq (30) we performed TNC-Cost, UCM, and GEM analyses. For the TNC-Cost analysis, we followed the procedure described in [20]. For the UCM analysis and GEM analysis, we computed the Jacobian of the action-to-outcome function at the mean action and the orthogonal direction.

4.5 6D throwing experimental protocol and data analysis

Individual release strategies were obtained from the experimental dataset acquired in our previous study [22], where twenty right-handed participants (10 females, 10 males; age: 28.2 ± 6.8 years) performed a series of overarm throws, starting from a fixed initial position. All participants signed an informed consent form in accordance with the Declaration of Helsinki. The data collection was carried out in accordance with Italian laws and European Union regulations on experiments involving human participants. The protocol was approved by the Ethical Review Board of the Santa Lucia Foundation (Prot. CE/PROG.542). Participants were instructed to hit one of four circular targets arranged on a vertical target board placed at 6 m from the initial position (marked with a sign on the floor) and to start from a fixed posture (standing with the arm along the body). The four targets were custom made and consisted in white circles of 40 cm diameter, arranged on a rectangular layout on the target board. The distances between the centers were 70 cm vertically and 80 cm horizontally, similarly to Fig 1. Moreover, the targets midpoints in the horizontal direction were shifted with respect to the projected initial position of the participant: the left and right targets were centered respectively at 60 cm to the left and 20 cm to the right of the projected initial position of the thrower's midline. An opto-electronic system (OptiTrack, NaturalPoint, Inc., Corvallis, OR, United States) operating at 120 Hz was used to capture whole-body kinematic information of the participants throwing actions and the corresponding ball trajectories.

For each trial and participant, the release action a , was obtained by fitting each of the three spatial components of the ball path with a 3rd-order polynomial function, and therefore the release position p_0 and velocity v_0 were obtained from the zero and first order coefficients, respectively. Then, this release action was used off-line in (27) and (25) to generate 'ideal' ball paths and scores, respectively, which were not influenced by friction and/or spinning effect of the ball. Trials in which the ball path did not intersect the target plane, or for which the ball was partially tracked by the optical system, were excluded from the analysis (13% of the total number of throws; we verified that differences in the number of trials across participants and conditions affect the values of some of the decomposition parameters only for sets with fewer trials than our smallest set). The error distribution, across trials, participants and target conditions, between experimental and ideal performance (mean squared error) is shown in Fig 9 (mean \pm SD: $0.0012 \pm 0.0912 m^2$). The dataset is available in S1 Dataset.

We next assessed the validity of the assumption that the individual action distribution is sufficiently localized around the mean action such that the second-order approximation of the sample mean score given in Eq (29) as the sum $\alpha + \beta$ is adequate. To do so we computed the

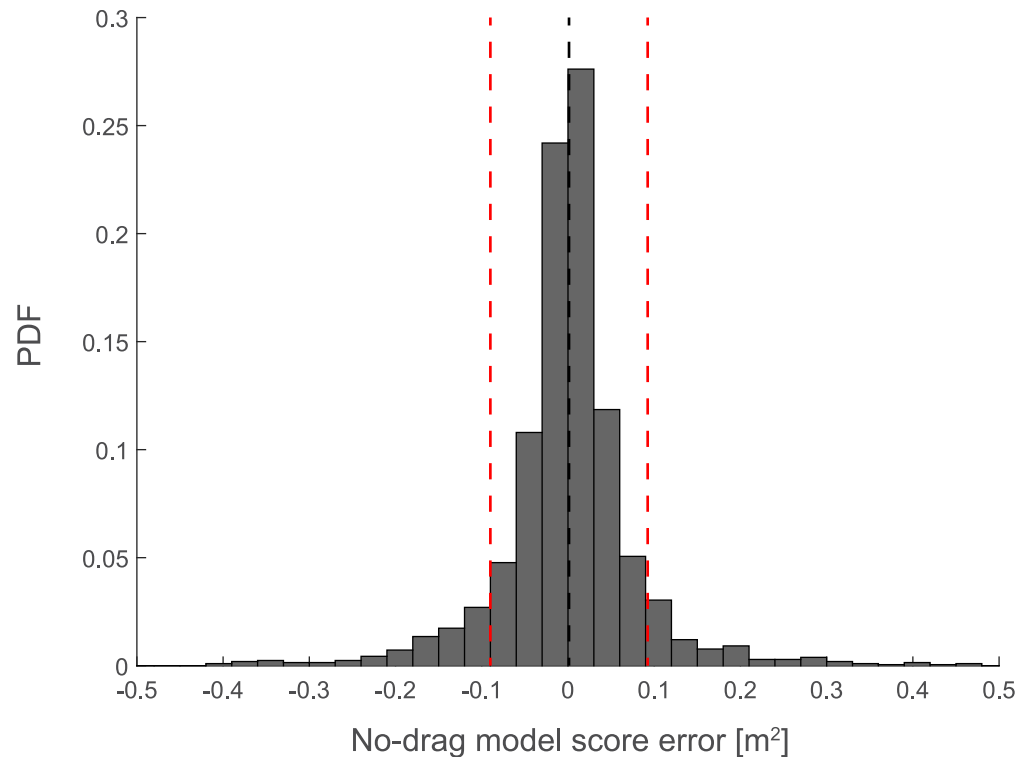


Fig 9. Error distribution between experimental score and score predicted with the no-drag model in Eq (29). Black and red vertical dashed lines indicate mean and $\pm 1SD$ of the distribution, respectively.

<https://doi.org/10.1371/journal.pone.0253626.g009>

fraction of variance accounted for (VAF) by the approximation, defined as:

$$VAF = \left[1 - \frac{\text{var}(E[\pi]_i - \alpha_i - \beta_i)}{\text{var}(E[\pi]_i)} \right] 100\% \tag{31}$$

VAF is therefore defined as the variance of the error between the individual performance (sample mean score) and the second order approximation normalized by the variance of the population performance.

In summary, for each participant and for each target \mathbf{x}_T , we estimated the mean-quadratic-error $E[\pi]$, the mean action $\bar{\mathbf{a}}$, and the action covariance Σ^a , with their respective sample mean and covariance: $E[\pi] = \frac{\sum_i s^x(\mathbf{x}_i(\mathbf{a}_i), \mathbf{x}_T)}{N}$, $\bar{\mathbf{a}} = \frac{\sum_i \mathbf{a}_i}{N}$, and $\Sigma^a = \frac{\sum_i (\mathbf{a}_i - \bar{\mathbf{a}})^T (\mathbf{a}_i - \bar{\mathbf{a}})^T}{N-1}$, where N is the total number of successful actions, or trials, executed for target \mathbf{x}_T .

4.6 Statistical analysis

To assess the effect of target and participant identity on the score and on the indicators extracted from the Hessian-based decomposition, we fit a linear mixed-effects model (Matlab function *fitlme*) with target and participants identity (categorical variables) as fixed effects and participants intercept as random effect

$$y \sim 1 + targetID + participantID + (1|participantID) \tag{32}$$

and we performed an analysis of variance for linear mixed-effects model (Matlab function *anova*), thus taking into account the repeated measures design.

Supporting information

S1 Appendix. (A.1) Relation to Muller & Sternad 2004; (A.2) Relation to Cohen & Sternad 2009 (B) Examples of individual distributions of release parameters in unconstrained 6D throwing (C) Score sensitivity (Hessian) and action variability (covariance) in unconstrained 6D throwing; (D) Coordinate invariance; (E) Non-zero Hessian's eigenvalues in the presence of redundant actions.

(PDF)

S1 Dataset. Throwing actions (release parameters), outcomes (ball arrival positions), and scores (squared distance of the arrival position from the center of the target) for each participant (P1-P20) and target (T1-T4) are included in a table (CSV file format). Each throw is a row. The table includes the following columns: Target ID, target positions (x, y, z), participant ID, release position and velocities ($x_0, y_0, z_0, v_{x_0}, v_{y_0}, v_{z_0}$), ball landing positions (x, z) computed with the no-drag model, score computed with the no-drag model, experimental score. (CSV)

Acknowledgments

We thank Etienne Burdet for providing many insightful comments and suggestions and Maura Mezzetti for helping with the statistical analysis.

Author Contributions

Conceptualization: Paolo Tommasino, Antonella Maselli, Andrea d'Avella.

Data curation: Paolo Tommasino, Antonella Maselli.

Formal analysis: Paolo Tommasino, Antonella Maselli, Andrea d'Avella.

Funding acquisition: Francesco Lacquaniti, Andrea d'Avella.

Investigation: Paolo Tommasino, Antonella Maselli, Andrea d'Avella.

Methodology: Paolo Tommasino, Domenico Campolo, Andrea d'Avella.

Software: Paolo Tommasino, Antonella Maselli, Andrea d'Avella.

Visualization: Paolo Tommasino, Antonella Maselli, Andrea d'Avella.

Writing – original draft: Paolo Tommasino, Antonella Maselli, Domenico Campolo, Andrea d'Avella.

Writing – review & editing: Paolo Tommasino, Antonella Maselli, Domenico Campolo, Francesco Lacquaniti, Andrea d'Avella.

References

1. Kording KP, Wolpert DM. The loss function of sensorimotor learning. *Proceedings of the National Academy of Sciences*. 2004; 101(26):9839–9842. <https://doi.org/10.1073/pnas.0308394101> PMID: 15210973
2. Hung YC, Kaminski TR, Fineman J, Monroe J, Gentile AM. Learning a multi-joint throwing task: a morphometric analysis of skill development. *Exp Brain Res*. 2008; 191(2):197–208. <https://doi.org/10.1007/s00221-008-1511-9> PMID: 18670769
3. Schmidt RA, Lee TD. *Motor control and learning: A behavioral emphasis*, 4th ed. Champaign, IL, US: Human Kinetics; 2005.
4. Dhawale AK, Smith MA, Ölveczky BP. The Role of Variability in Motor Learning. *Annu Rev Neurosci*. 2017; 40(1):479–498. <https://doi.org/10.1146/annurev-neuro-072116-031548> PMID: 28489490

5. Harris CM, Wolpert DM. Signal-dependent noise determines motor planning. *Nature*. 1998; 394(6695):780–784. <https://doi.org/10.1038/29528> PMID: 9723616
6. Osborne LC, Lisberger SG, Bialek W. A sensory source for motor variation. *Nature*. 2005; 437(7057):412–416. <https://doi.org/10.1038/nature03961> PMID: 16163357
7. 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*. 2004; 30(1):212–233. <https://doi.org/10.1037/0096-1523.30.1.212> PMID: 14769078
8. Bernstein NA. *Dexterity and Its Development*. Psychology Press; 1996.
9. Fowler CA, Turvey MT. 1—Skill Acquisition: An Event Approach with Special Reference to Searching for the Optimum of a Function of Several Variables. In: Stelmach GE, editor. *Information Processing in Motor Control and Learning*. Academic Press; 1978. p. 1–40. Available from: <http://www.sciencedirect.com/science/article/pii/B9780126659603500062>.
10. Newell KM. Motor Skill Acquisition. *Annu Rev Psychol*. 1991; 42(1):213–237. <https://doi.org/10.1146/annurev.ps.42.020191.001241> PMID: 2018394
11. Sternad D. It's not (only) the mean that matters: variability, noise and exploration in skill learning. *Current Opinion in Behavioral Sciences*. 2018; 20:183–195. <https://doi.org/10.1016/j.cobeha.2018.01.004> PMID: 30035207
12. Haar S, van Assel CM, Faisal AA. Motor learning in real-world pool billiards. *Sci Rep*. 2020; 10(1):20046. <https://doi.org/10.1038/s41598-020-76805-9> PMID: 33208785
13. Sternad D, Abe MO, Hu X, Müller H. Neuromotor Noise, Error Tolerance and Velocity-Dependent Costs in Skilled Performance. *PLoS Comput Biol*. 2011; 7(9):e1002159. <https://doi.org/10.1371/journal.pcbi.1002159> PMID: 21966262
14. Thorp EB, Kording KP, Mussa-Ivaldi FA. Using noise to shape motor learning. *Journal of Neurophysiology*. 2017; 117(2):728–737. <https://doi.org/10.1152/jn.00493.2016> PMID: 27881721
15. Scholz JP, Schöner G. The uncontrolled manifold concept: identifying control variables for a functional task. *Experimental Brain Research*. 1999; 126(3):289–306. <https://doi.org/10.1007/s002210050738> PMID: 10382616
16. Latash ML, Scholz JP, Schöner G. Motor Control Strategies Revealed in the Structure of Motor Variability. *Exercise and Sport Sciences Reviews*. 2002; 30(1):26–31. <https://doi.org/10.1097/00003677-200201000-00006> PMID: 11800496
17. Cusumano JP, Cesari P. Body-goal Variability Mapping in an Aiming Task. *Biol Cybern*. 2006; 94(5):367–379. <https://doi.org/10.1007/s00422-006-0052-1> PMID: 16501988
18. Cusumano JP, Dingwell JB. Movement variability near goal equivalent manifolds: Fluctuations, control, and model-based analysis. *Human Movement Science*. 2013; 32(5):899–923. <https://doi.org/10.1016/j.humov.2013.07.019> PMID: 24210574
19. Yang JF, Scholz JP. Learning a throwing task is associated with differential changes in the use of motor abundance. *Exp Brain Res*. 2005; 163(2):137–158. <https://doi.org/10.1007/s00221-004-2149-x> PMID: 15657698
20. Cohen RG, Sternad D. Variability in motor learning: relocating, channeling and reducing noise. *Exp Brain Res*. 2009; 193(1):69–83. <https://doi.org/10.1007/s00221-008-1596-1> PMID: 18953531
21. Zhang Z, Sternad D. Back to reality: differences in learning strategy in a simplified virtual and a real throwing task. *J Neurophysiol*. 2021; 125(1):43–62. <https://doi.org/10.1152/jn.00197.2020> PMID: 33146063
22. Maselli A, Dhawan A, Cesqui B, Russo M, Lacquaniti F, d'Avella A. Where Are You Throwing the Ball? I Better Watch Your Body, Not Just Your Arm! *Front Hum Neurosci*. 2017; 11:505. <https://doi.org/10.3389/fnhum.2017.00505> PMID: 29163094
23. Maselli A, Dhawan A, Russo M, Cesqui B, Lacquaniti F, d'Avella A. A whole body characterization of individual strategies, gender differences, and common styles in overarm throwing. *J Neurophysiol*. 2019; 122(6):2486–2503. <https://doi.org/10.1152/jn.00011.2019> PMID: 31577474
24. Venkadesan M, Mahadevan L. Optimal strategies for throwing accurately. *R Soc open sci*. 2017; 4(4):170136. <https://doi.org/10.1098/rsos.170136> PMID: 28484641
25. Campolo D, Widjaja F, Xu H, Ang WT, Burdet E. Analysis of accuracy in pointing with redundant hand-held tools: a geometric approach to the uncontrolled manifold method. *PLoS Comput Biol*. 2013; 9(4):e1002978. <https://doi.org/10.1371/journal.pcbi.1002978> PMID: 23592956
26. Smeets JBJ, Louw S. The contribution of covariation to skill improvement is an ambiguous measure: Comment on Müller and Sternad (2004). *Journal of Experimental Psychology: Human Perception and Performance*. 2007; 33(1):246–249. <https://doi.org/10.1037/0096-1523.33.1.246> PMID: 17311491

27. Abe MO, Sternad D. Directionality in distribution and temporal structure of variability in skill acquisition. *Front Hum Neurosci*. 2013; 7:225. <https://doi.org/10.3389/fnhum.2013.00225> PMID: 23761742
28. Sternad D, Park SW, Müller H, Hogan N. Coordinate Dependence of Variability Analysis. *PLoS Comput Biol*. 2010; 6(4):e1000751. <https://doi.org/10.1371/journal.pcbi.1000751> PMID: 20421930
29. Müller H, Frank TD, Sternad D. Variability, covariation, and invariance with respect to coordinate systems in motor control: Reply to Smeets and Louw (2007). *Journal of Experimental Psychology: Human Perception and Performance*. 2007; 33(1):250–255. <https://doi.org/10.1037/0096-1523.33.1.250> PMID: 17311492
30. John J, Dingwell JB, Cusumano JP. Error Correction and the Structure of Inter-Trial Fluctuations in a Redundant Movement Task. *PLoS Comput Biol*. 2016; 12(9):e1005118. <https://doi.org/10.1371/journal.pcbi.1005118> PMID: 27643895
31. Shadmehr R, Orban de Xivry JJ, Xu-Wilson M, Shih TY. Temporal Discounting of Reward and the Cost of Time in Motor Control. *Journal of Neuroscience*. 2010; 30(31):10507–10516. <https://doi.org/10.1523/JNEUROSCI.1343-10.2010> PMID: 20685993
32. Haith AM, Krakauer JW. Model-Based and Model-Free Mechanisms of Human Motor Learning. In: Richardson MJ, Riley MA, Shockley K, editors. *Progress in Motor Control*. vol. 782. New York, NY: Springer New York; 2013. p. 1–21. Available from: http://link.springer.com/10.1007/978-1-4614-5465-6_1.
33. Chen X, Holland P, Galea JM. The effects of reward and punishment on motor skill learning. *Current Opinion in Behavioral Sciences*. 2018; 20:83–88. <https://doi.org/10.1016/j.cobeha.2017.11.011>
34. Harding EF. 10. Quadratic Forms in Random Variables: Theory and Applications. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*. 1993; 156(2):326–327. <https://doi.org/10.2307/2982749>