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

Paolo Tommasino[1*], Antonella Maselli[1], Domenico Campolo[2], Francesco Lacquaniti[1,3], Andrea d'Avella[1,4*]

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

* corresponding authors
tommasinopaolo@gmail.com (PT), a.davella@hsantalucia.it (AdA)

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 the noise sensitivity of the action distribution. 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 overall noise magnitude, overall 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 suboptimal mean action for better tolerance-variability and higher action variability for better alignment with more tolerant directions in action space.

Author summary

Why do people differ in their performance of complex motor skills? In many real-life motor tasks achieving a goal requires selecting an appropriate high-dimensional action out of infinitely many goal-equivalent actions. Because of sensorimotor noise, we are unable to execute the exact same action twice and our performance depends on how accurate we are on average. Thus, to understand why people perform differently we need to characterize how their action distribution relates to their mean task score. While better performance is often associated to smaller variability around a more accurate mean action, performance also depends on the relationship between the directions of highest variability in action space and the directions in which action variability affects the most the outcome of the action. However, characterizing such geometric relationship when actions are high dimensional is challenging. In this work we introduce a method that allows to characterize the key performance-related features of the distribution of high-dimensional actions by a small set of indicators. We can then compare such indicators in different people performing a complex task (such as unconstrained throwing) and directly characterize the most skilled ones but also identify different strategies that distinguish people with similar performance.

1 Introduction

In many goal-directed human behaviors, such as throwing a projectile towards the center of a target, performance depends on how accurate the outcome of repeated actions is. 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 [16, 12]. In this perspective, the goal of a thrower would be that of minimizing/maximizing the mean score over repeated trials [30, 7]. Because of bias and noise in the sensorimotor transformations mapping goals into actions [11, 23], 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 [16, 21].

In many experimental as well as naturalistic scenarios, the relationship between motor actions and their outcomes is *redundant* [1] 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 (*actions*). 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 (they all result into 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, 22, 16, 27].

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 [28, 31]. In this perspective, the relationship between action distribution and outcome variability in goal-directed behaviors has been addressed by computational approaches that take into account the geometry mapping between actions and outcomes near the solution manifold. Methods such as the Uncontrolled Manifold (UCM) [24, 17] and the Goal-Equivalent Manifold (GEM)[5, 6] typically approximate the (non-linear) solution manifold 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 [24, 17, 33] has shown that *covariation* between redundant actions is an important mechanism used by skilled performers to push "bad" motor variability along the solution manifold, 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 actions, i.e. factors depending on the local geometry of the goal function, rather than different amounts of action variability only.

The impact of the interplay between motor variability and task geometry on performance in goal-directed behaviors has been also investigated with an approach that needs no assumption on the smoothness of the action-to-score mapping, and relies on the use of surrogate data computations, rather than analytic descriptions, to explore its geometry [21, 4, 27]. The Tolerance-Noise-Covariation (TNC) method [21] 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 [4]) 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 computation of covariation cost depends on a numerical procedure that becomes cumbersome for high dimensional action spaces [27], the TNC 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 [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 [19, 18], 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 and 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 a free 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, 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 Methods

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, and the covariance of the action distribution. Fig 2A illustrates the variables and functions describing the relationship between an action ($\boldsymbol{a}$) and a score (or loss, π) assigned to the outcome ($\boldsymbol{x}$) of the action. Such score represents a scalar measure of inaccuracy of the outcome with respect to the goal (or target, $\boldsymbol{x}_T$) and, thus, it is a composite function of the action-to-outcome mapping ($\boldsymbol{x} = f(\boldsymbol{a})$) and the outcome-to-score mapping ($\pi = s^a(\boldsymbol{x}; \boldsymbol{x}_T)$). Fig 2B provides a graphical illustration of the

decomposition approach in the case of two-dimensional actions. 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 ($\bar{\boldsymbol{a}}$) and covariance (Σ) while the local geometry of the action-to-score mapping by the Hessian (H) computed at the mean action. We will show in the next section how the expected value of the score ($E(\pi)$), when defined as a loss function, can be decomposed as the sum of the score of the mean action ($\alpha = s^a(\bar{\boldsymbol{a}})$) and a *tolerance-variability* index (β) expressed as the product of three terms: sensitivity of the action-to-score mapping to action variability (reciprocal of the *tolerance* τ), due to the magnitude of the Hessian eigenvalues; amount of variability or *uncorrelated noise* (η), due to the magnitude of the covariance eigenvalues; *alignment* (θ) between the eigenvectors of the Hessian and the eigenvectors of the covariance.

2.1 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 $\boldsymbol{a}_t \in A \subset R^n$ with some random noise such that the action distribution can be described by a probability density function (p.d.f.) $p_{\boldsymbol{A}}$. Let us also assume that, at every trial, the action receives a score point $\pi_t = s^a(\boldsymbol{a}_t)$ through a *score* (or loss) function $s^a : \boldsymbol{a} \in R^n \rightarrow R^+$ 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 $\boldsymbol{x}(\boldsymbol{a})$ and the target position $\boldsymbol{x}_T$. Following a decision theory approach, we define the *performance* of a motor strategy $p_{\boldsymbol{A}}$ as the *expected score*:

$$E[\pi] = \int_{A \subset R^n} s^a(a_1, a_2, \ldots, a_n) p_{\boldsymbol{A}}(a_1, a_2, \ldots a_n) da_1 da_2 \ldots da_n. \tag{1}$$

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 $p_{\boldsymbol{A}}$ 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 real practice, the score function may be highly non-linear and the action space high dimensional ($n >> 1$) making difficult to find an analytic solution to (1). Therefore, we do not usually solve (1) 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.

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 (1).

Let's assume that an individual selects motor actions according to a p.d.f. with *expected* or *mean action* $E[\boldsymbol{a}] = \bar{\boldsymbol{a}} \in R^n$, and covariance matrix $\Sigma^a \in P^n$ (symmetric and positive definite). In other words, at every trial t, an individual generates an action $\boldsymbol{a}_t$ according to the following stochastic model:

$$\boldsymbol{a}_t = \bar{\boldsymbol{a}} + \boldsymbol{\delta_a} \tag{2}$$

where $\boldsymbol{\delta_a} \in R^n$ is the *stochastic component* of the individual strategy, which, at every trial, 'perturbs' the mean action $\bar{\boldsymbol{a}}$. 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 $\bar{\boldsymbol{a}}$, the score $s^a(\boldsymbol{a}_t)$ of an action $\boldsymbol{a}_t$, can be approximated with the following second-order Taylor expansions:

$$s^a(\bar{\boldsymbol{a}} + \boldsymbol{\delta_a}) \approx s^a(\bar{\boldsymbol{a}}) + [\boldsymbol{\nabla_a} s^s(\bar{\boldsymbol{a}})]^T \boldsymbol{\delta_a} + \frac{1}{2}\boldsymbol{\delta_a}^T H(\bar{\boldsymbol{a}})\boldsymbol{\delta_a} \tag{3}$$

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

$$H_{ij}(\bar{\boldsymbol{a}}) = \frac{\partial^2 s^a}{\partial a_i \partial a_j}(\bar{\boldsymbol{a}}). \tag{4}$$

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

Inserting (3) in the integral (1), we can write the expected action score as the sum of three terms:

$$E[\pi] \approx E[s^a(\bar{\boldsymbol{a}})] + [\boldsymbol{\nabla_a} s^s(\bar{\boldsymbol{a}})]^T E[\boldsymbol{\delta_a}] + E[\frac{1}{2}\boldsymbol{\delta_a}^T H(\bar{\boldsymbol{a}})\boldsymbol{\delta_a}] \tag{5}$$

the first term $E[s^a(\bar{\boldsymbol{a}})]$ is simply *the score of the average action* $s^a(\bar{\boldsymbol{a}})$ given that the expected value of a constant is the constant itself. The second term, $[\boldsymbol{\nabla_a} s^s(\bar{\boldsymbol{a}})]^T E[\boldsymbol{\delta_a}]$, vanishes whenever the quadratic approximation is evaluated at the mean action $\bar{\boldsymbol{a}}$, 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 $trace(\frac{1}{2}H(\bar{\boldsymbol{a}})\Sigma^a)$ [10]. 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 (1) can be (locally) approximated as:

$$E[\pi] \approx s^a(\bar{\boldsymbol{a}}) + trace(\frac{1}{2}H(\bar{\boldsymbol{a}})\Sigma^a) = \alpha(\bar{\boldsymbol{a}}) + \beta(\bar{\boldsymbol{a}}, \Sigma^a) \tag{6}$$

where:

$$\alpha(\bar{\boldsymbol{a}}) \equiv s^a(\bar{\boldsymbol{a}}) \tag{7}$$

is *the score of the average action*, and

$$\beta(\bar{\boldsymbol{a}}, \Sigma^a) \equiv trace\left(\frac{1}{2}H(\bar{\boldsymbol{a}})\Sigma^a\right) \tag{8}$$

is the *tolerance-variability index* which captures how local non-linearities of the action score ($H(\bar{\boldsymbol{a}})$) and variability in the action strategy (Σ^a) affect (increase/decrease) performance, i.e. the mean score $E[\pi]$. Notice 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 uncorrelated noise*, *the score tolerance* and the *alignment*.

2.2 Hessian-based decomposition of motor skills

Principal variability directions and *uncorrelated 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):

$$\Sigma^a = U^{\Sigma}\Lambda^{\Sigma}U^{\Sigma^T} = \left[\ \boldsymbol{u}_1^{\Sigma}, \boldsymbol{u}_2^{\Sigma}, ..., \boldsymbol{u}_n^{\Sigma}\ \right] diag(\lambda_1^{\Sigma}, \lambda_2^{\Sigma}, \ldots, \lambda_n^{\Sigma}) \left[\ \boldsymbol{u}_1^{\Sigma}, \boldsymbol{u}_2^{\Sigma}, ..., \boldsymbol{u}_n^{\Sigma}\ \right]^T. \tag{9}$$

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

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

$$\eta(\Sigma^a) \equiv trace(\Sigma^a) = \sum_{j=1}^{n} \sigma_{jj} = \sum_{l=1}^{n} \lambda_l^{\Sigma} \tag{10}$$

Principal sensitivity directions and *tolerance* τ. When locally, i.e. around the mean action $\bar{\boldsymbol{a}}$, 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:

$$H = U^H\Lambda^H U^{H^T} = \left[\ \boldsymbol{u}_1^h, \boldsymbol{u}_2^h, ..., \boldsymbol{u}_n^h\ \right] diag(\lambda_1^H, \lambda_2^H, \ldots, \lambda_n^H) \left[\ \boldsymbol{u}_1^h, \boldsymbol{u}_2^h, ..., \boldsymbol{u}_n^h\ \right]^T. \tag{11}$$

The diagonal matrix $\Lambda^H = diag(\lambda_1^H, \lambda_2^H, \ldots, \lambda_n^H)$ contains the n singular values $\lambda_1^h, \lambda_2^h, ..., \lambda_n^h$ (with $\lambda_1^h \geq \lambda_2^h \geq, ..., \lambda_n^h$) of the Hessian matrix. The orthonormal matrix U^H contains the associated singular vectors $\left[\ \boldsymbol{u}_1^h, \boldsymbol{u}_2^h, ..., \boldsymbol{u}_n^h\ \right]$. 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* $\boldsymbol{u}_j^h$. 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 $\boldsymbol{\delta_a}$ which are directed along the i-th principal sensitivity vector $\boldsymbol{u}_i^H$. As a local, scalar measure, of 'total curvature' of the score, we define the *sensitivity* of the score as the $trace(H) = \sum_{i=1}^{n} h_{ii} = \sum_{l=1}^{n} \lambda_l^H$.

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

$$\tau(\bar{\boldsymbol{a}}) \equiv \frac{1}{trace\,(H(\bar{\boldsymbol{a}}))}. \tag{12}$$

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:

$$\beta = trace\left(\frac{1}{2}\frac{H}{trace(H)}\frac{\Sigma}{trace(\Sigma)}\right) trace(H) trace(\Sigma) = trace\left(\frac{1}{2}\bar{H}\bar{\Sigma}\right)\frac{\eta}{\tau} = \frac{\theta}{\tau}\eta \tag{13}$$

where the *alignment* θ:

$$\theta(\bar{\boldsymbol{a}}, \Sigma^a) = trace(\frac{1}{2}\bar{H}\bar{\Sigma}) \tag{14}$$

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 $\boldsymbol{u}_i^{\Sigma}$ are aligned with the directions of maximal sensitivity $\boldsymbol{u}_i^{H}$ 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:

$$E[\pi] \approx \alpha(\bar{\boldsymbol{a}}) + \beta(\bar{\boldsymbol{a}}, \Sigma^a) = \alpha(\bar{\boldsymbol{a}}) + \frac{\theta(\bar{\boldsymbol{a}}, \Sigma^a)}{\tau(\bar{\boldsymbol{a}})}\eta(\Sigma^a) \tag{15}$$

Score relevant subspace and variability in redundant motor tasks

For redundant tasks, the map $\boldsymbol{f}$, between actions and outcomes, i.e. $\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{a})$, maps the n-dimensional action space onto the m-dimensional task space (or outcome space), with $m < n$. In such case, the solution manifold will be an $n-m$-dimensional surface embedded into the n-dimensional action space, and the Hessian matrix, along the solution manifold will only have m non-zero eigenvalues (see S1 Appendix C), Therefore, the first m principal sensitivity directions $\boldsymbol{u}_i$ with $i = 1, 2, \ldots m$ of the action score function span a lower-dimensional subspace (*score-relevant subspace*) of the action space, and the Hessian matrix can be approximated with an $n \times n$ matrix of rank m:

$$H_m = U_m^H \Lambda_m^H U_m^{H^T} \quad \text{with } 1 \leq m < n \tag{16}$$

where U_m^H is an $n \times m$ matrix and Λ_m^H is an $m \times m$ diagonal matrix having the m-largest singular values on the diagonal.

For symmetric and positive-definite Hessian matrices, β can also be written as:

$$\beta = \text{trace}\left(\frac{1}{2}H^{\frac{1}{2}}\Sigma H^{\frac{1}{2}}\right) \tag{17}$$

with $H^{\frac{1}{2}} = U^H \Lambda^{H^{\frac{1}{2}}} U^{H^T}$. Let m be the number of score-relevant dimensions near the solution manifold, then β can also be approximated as:

$$\beta \approx \text{trace}\left(\frac{1}{2}\left(U_m^H \Lambda_m^{H^{\frac{1}{2}}} U_m^{H^T}\right)\Sigma\left(U_m^H \Lambda_m^{H^{\frac{1}{2}}} U_m^{H^T}\right)\right) \tag{18}$$

The internal product $U_m^{H^T}\Sigma U_m^H$ represents the *score-relevant* variability Σ_H:

$$\Sigma_H = U_m^{H^T}\Sigma U_m^H = E\left[\left(U_m^H\boldsymbol{\delta_a}\right)^T\left(U_m^H\boldsymbol{\delta_a}\right)\right] \tag{19}$$

i.e., the fraction of action variability that is tangent to the score-relevant dimensions. Near the solution manifold the β index is therefore the trace of the score-relevant variability further *amplified/attenuated* by the local sensitivity of the score function:

$$\beta \approx \text{trace}\left(\frac{1}{2}U_m^H\underbrace{\Lambda_m^{H^{\frac{1}{2}}}\Sigma_H\Lambda_m^{H^{\frac{1}{2}}}}_{B}U_m^{H^T}\right) = \frac{1}{2}\text{trace}(B) \tag{20}$$

given that the columns of U_m^H are unitary vectors.

Lastly, notice that in the presence of motor redundancy, the distribution of motor actions can be highly concentrated in a lower dimensional space of the action space. In such cases, the covariance matrix becomes almost singular and only $k < n$ singular values are significantly different from zero. Then, the first k principal variability directions $\boldsymbol{u}_i^\Sigma, i = 1\ldots k$ span a k-dimensional subspace of the action space and the covariance can be approximated with an $n \times n$ matrix of rank k:

$$\Sigma_k^a = U_k^\Sigma\Lambda_k^\Sigma U_k^{\Sigma^T} \tag{21}$$

where U_k^Σ is an $n \times k$ matrix and Λ_k^Σ a $k \times k$ diagonal matrix having the k-largest singular values on the diagonal. The columns of the matrix $U_k^\Sigma\Lambda_k^\Sigma$, define the first k principal components of the action variability.

2.3 Application to throwing tasks

Task score vs. action score

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 [19] 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 [16, 25], we assume that participants optimize an accuracy score which penalizes the squared error between the outcome $\boldsymbol{x}$ of a release action and the target position $\boldsymbol{x}_T$ [25]:

$$\pi = s^x(\boldsymbol{x}, \boldsymbol{x}_T) = (\boldsymbol{x} - \boldsymbol{x}_T)^T(\boldsymbol{x} - \boldsymbol{x}_T) \equiv e^2. \tag{22}$$

Written in this form, the accuracy score, represents a *task score* $s^x : R^2 \to R^+$ which penalizes the two-dimensional action outcomes $\boldsymbol{x}$ with a scalar score π. To find the relationship between

release actions and quadratic (task) error, i.e. the *action-to-score function* $s^a : \boldsymbol{a} \in R^6 \rightarrow R^+$, we need to express the task score (22) as a function of the release parameters.

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

$$\boldsymbol{f}(\boldsymbol{p}_0, \boldsymbol{v}_0, t) = \begin{cases} x(t) = x_0 + v_{x0} t \\ y(t) = y_0 + v_{y0} t \\ z(t) = z_0 + v_{z0} t - \frac{1}{2} g t^2 \end{cases} \tag{23}$$

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

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 (23)). Hence, at the time of impact, the projectile will hit the board at:

$$\boldsymbol{x}_i = \boldsymbol{f}(\boldsymbol{a}, t_I) = \begin{cases} x_0 + v_{x0} \frac{y_b - y_0}{v_{y0}} \\ y_b \\ z_0 + v_{z0} \frac{y_b - y_0}{v_{y0}} - \frac{1}{2} g \left(\frac{y_b - y_0}{v_{y0}} \right)^2 \end{cases} \tag{24}$$

Substituting the above system of equations into (22), let us writing the *action score* as:

$$e^2 = s^x(\boldsymbol{x}_i, \boldsymbol{x}_T) = s^x(\boldsymbol{f}(\boldsymbol{a}, t_I), \boldsymbol{x}_T) = s^a(\boldsymbol{a}, \boldsymbol{x}_T) \tag{25}$$

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

$$e^2 = s^a(\boldsymbol{a}, \boldsymbol{x}_T) = \left(x_0 + v_{x0} \frac{y_b - y_0}{v_{y0}} - x_T \right)^2 + \left(z_0 + v_{z0} \frac{y_b - y_0}{v_{y0}} - \frac{1}{2} g \left(\frac{y_b - y_0}{v_{y0}} \right)^2 - z_T \right)^2 \tag{26}$$

Given an individual release strategy with mean action $\bar{\boldsymbol{a}}_0$ and covariance Σ^a and aiming at hitting a desired target $\boldsymbol{x}_T$, the mean squared error ($E[e^2] = E[\pi]$) can be approximated with (6):

$$E[\pi] = E[e^2] \approx \alpha(\bar{\boldsymbol{a}}) + \beta(\bar{\boldsymbol{a}}, \Sigma^a) = \alpha(\bar{\boldsymbol{a}}) + \frac{\theta(\bar{\boldsymbol{a}}, \Sigma^a)}{\tau(\bar{\boldsymbol{a}})} \eta(\Sigma^a) \tag{27}$$

where $\alpha(\bar{\boldsymbol{a}})$ is just (26) evaluated at $\bar{\boldsymbol{a}}$, 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 (26) evaluated at $\bar{\boldsymbol{a}}$.

In this work, the 6×6 Hessian matrix of (26) is calculated with the MATLAB Symbolic Toolbox for each target condition and for each individual strategy.

A simulated 2D example

Equation (27) tells us that the mean squared error of a release strategy, hence the performance of a thrower, depends on three factors: the *mean action* $\bar{\boldsymbol{a}}$, the *action variability* Σ^a and the local geometry of the score $H(\bar{\boldsymbol{a}})$, which increase/decrease the β score by *amplifying/attenuating* stochastic action perturbations $\boldsymbol{\delta_a}$. Clearly, an optimal or close-to optimal release strategy, is the one that on average gets zero or close-to zero penalty score, therefore, skilled throwers should have α and β parameters close to zero. To give a visual representation of our approach Fig 3A shows a toy model of a dart throwing task. In this task, the action score (26) is assumed to be dependent only on the longitudinal and vertical release velocity, while (22) penalizes only the vertical errors with respect to target position. The score of an action as well as the score of its task outcome, is represented with a gray-scale color code: light/dark colors, represents actions which receives low/high penalty points, respectively. The yellow line represents the *solution manifold*, i.e. the set of *optimal actions* which results into 0 penalty score. The local tolerance of the score, onto and close-to the solution manifold, is represented with the red ellipses, whose major and minor axes are, the first and second principal sensitivity directions of the score, respectively. Hence, the longer the axis the more sensitive/less tolerant the score is to perturbations directed along the axis direction. Notice that the ellipses are anisotropic, with $\lambda_1^H >> \lambda_2^H$, suggesting that, near the solution manifold, there is only one direction in action space which *is relevant* for the score, and its expected value. This direction however, is not constant, but changes along the solution manifold, as shown by the different orientations of the ellipses. Also notice that, the smaller the ratio $\frac{v_{z0}}{v_{y0}}$, the smaller the curvature (smaller ellipses), hence the larger the noise tolerance. In other words, the score is more tolerant when the average longitudinal velocity of a release strategy is larger than the average vertical velocity. Fig 3B shows three simulated individual strategies, each drawing actions according to a bi-dimensional Gaussian distribution with mean action $\bar{\boldsymbol{a}}_i$ and covariance Σ_i, with $i = g, b, p$ (g = green, b = blue, p = purple), as reported in Table 1. Each colored circle in 3B *left* represents an action drawn from each of the distributions and Fig 3B *center* and *right* show the ball trajectories of each action and the score distribution of each strategy, respectively. Fig 3C shows the decomposition of the mean score of each strategy according to the performance-analysis presented in previous section. Since $\bar{\boldsymbol{a}}_g$ does not belong to the solution manifold, $\bar{\boldsymbol{a}}_g$ is not optimal and hence its score is not zero, i.e. $\alpha \neq 0$. This results into 'biased' throwing outcomes, whose average landing position does not coincide with the desired target position, as also shown by the green distribution in Fig 3B *center*. Notice, however, that while the g strategy has the worse α, it has the best (lowest) β compared to b and m. The β score depends on the interplay between action variability and local curvature and hence on the three parameters η, τ and θ shown in Fig 3C. Also notice that, in this example, the three strategies have been chosen to have the same level of noise $\eta_b = trace(\Sigma_b) = \eta_g = trace(\Sigma_g) = \eta_p = trace(\Sigma_p)$, hence differences in β are only due to the local tolerance τ and the alignment θ between principal sensitivity directions and principal variability directions. The blue and purple strategies, have a similar alignment θ, given that, for both strategies, the principal component (direction of greatest variability) is almost parallel to the direction of maximal sensitivity. However, compared to the purple strategy, the blue strategy can afford a much lower β because its mean action is located in a more tolerant region of the score (as shown by the wider white area), i.e. $\tau_b >> \tau_p$. While the

	$\bar{\boldsymbol{a}}[m/s]$	$\Sigma^a\ [m/s]^2$
B	$\begin{bmatrix} 10 & 2.95 \end{bmatrix}$	$10^{-3} \cdot \begin{bmatrix} 8.4 & 11 \\ 11 & 21 \end{bmatrix}$
G	$\begin{bmatrix} 8 & 3.20 \end{bmatrix}$	$10^{-3} \cdot \begin{bmatrix} 8 & -10 \\ -10 & 20 \end{bmatrix}$
P	$\begin{bmatrix} 6 & 4.91 \end{bmatrix}$	$10^{-3} \cdot \begin{bmatrix} 8.4 & 11 \\ 11 & 21 \end{bmatrix}$

Table 1. Mean and covariance matrices used to simulate the Gaussian distribution of the blue, the green and the purple strategy of the example in Fig.3.

comparison between the blue and purple strategy emphasizes how the mean action can impact tolerance and hence the average score, the comparison between the green and the blue strategy emphasizes the effect of the alignment: the green strategy has a smaller β, despite $\tau_g < \tau_b$, because the direction of maximal action variability $\boldsymbol{u}_1^{\Sigma}$ *is less aligned*, compared to the blue strategy, with the local direction of maximal sensitivity $\boldsymbol{u}_1^{H}$.

2.4 Experimental protocol and data analysis

Individual release strategies were obtained from the experimental dataset acquired in our previous study [19], 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 throwers 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 $\boldsymbol{a}$, 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 $\boldsymbol{p}_0$ and velocity $\boldsymbol{v}_0$ were obtained from the zero and first order coefficients, respectively. Then, this release action was used off-line in (24) and (22) 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 4 (mean $\pm$ SD: $0.0012 \pm 0.0912 m^2$). The dataset is available in S2 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 equation (26) as the sum $\alpha + \beta$ is adequate. To do so we computed the fraction of variance accounted for (VAF) by the approximation, defined as:

$$\text{VAF} = \left[1 - \frac{\text{var}(E[\pi]_i - \alpha_i - \beta_i)}{\text{var}(E[\pi]_i)}\right] 100\% \tag{28}$$

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

3 Results

3.1 Examples of individual distributions of throwing parameters

When actions are high-dimensional, as in unconstrained throwing, it is impossible to visualize both the score and the individual strategies in a single plot. Fig 5 shows the action score and individual throwing strategies in terms of pairs of release parameters (9 of the 15 possible pairs of 3 position and 3 velocity variables; *rows*) for five exemplary participants (*columns*) throwing at target T1. Participants have been sorted from left to right according to their average release speed, hence $P10$ is the slowest thrower and $P11$ the fastest. The gray-scale shaded areas indicate the score associated to each pair of release parameters, i.e. the squared distance between the ball arrival position on the vertical target plane and the center of the target (see Eq 26). The domain of each position and velocity variable corresponds to the population mean $\pm$ 3 standard deviations, while all the remaining release parameters are considered constant and fixed to the subject-specific mean action. The plots have 15 different gray-scale levels: the white area defines actions which have a score smaller or equal to 0.04 m^2, i.e. actions that land inside the target, which has radius 0.2 m. Notice 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, the more tolerant is the action-to-score function to stochastic perturbations. For instance, in the $v_{x_0} - v_{z_0}$ plane (*8-th row*), the lowest penalty (white) area looks like an ellipse, whose orientation suggests that the v_{z_0} direction is less tolerant/more sensitive to action variability and whose size, increasing with release speed, suggests that tolerance is higher for higher speed. In general, position variables appear more tolerant than velocity variables and the

shapes of the white regions in the $p_{y_0} - p_{z_0}$ plane (*third row*), in the same-axis position-velocity planes (*fourth to sixth rows*), and in the $v_{y_0} - v_{z_0}$ plane (*ninth row*) indicate that release variables corresponding to the lowest score (throws hitting the target) are negatively correlated.

The distributions of the release parameters (*red circles*), summarized in each plot of Fig 5 in terms of mean (*blue circle*) and covariance (two-standard deviations, *blue ellipse*), also differs remarkably across planes and participants. For instance, $P11$ (*fifth column*), on average, releases the ball with a faster vertical velocity (v_{z_0}), compared to other participants, such as $P10$ (*first column*) and $P1$ (*third column*), who instead throw with higher longitudinal velocity (v_{y_0}). The amount of variability in the $v_{x_0} - v_{z_0}$ (*eight row*) and $v_{y_0} - v_{z_0}$ (*ninth row*) planes also differs between the same three participants, as $P11$ shows a much wider covariance ellipses than $P1$ who, in turn, shows a wider covariance ellipses than $P10$. However, interestingly, their performance is similar (see Fig 1), as the large variability of $P11$ in the $v_{x_0} - v_{z_0}$ and $v_{y_0} - v_{z_0}$ planes is partially compensated by being in a more tolerant region of the action score.

Also to notice are different patterns of covariance/correlation across individual strategies. For instance, $P10$, who is the least variable participant, shows no correlation between v_{y_0} and v_{z_0}, while $P11$, who is the most variable participant, shows a negative correlation: reducing the vertical release velocity proportionally to an increase in the longitudinal release velocity. This allows $P11$ to remain in the lowest penalty region as much as $P1$ and $P10$ and hence to have a similar mean score. Notice that $P15$ has a mean release velocity located close to the edge of the white region and hence on average does not hit the center of the target. However, $P15$ has a mean score similar to those of $P10$, $P1$, and $P11$, in part due to the small and, in most cases, well aligned covariance ellipses. Finally, $P18$, the best performing participant, shows mean release parameters at the center of all white regions and small and well aligned covariance ellipses.

In sum, the examination of the distribution of several pairs of release action parameters and the associated score suggests that individual throwing strategies differs in terms of mean action, action variability, and relationship between action variability and geometry of the action-to-score function. 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 introducing the Hessian-based decomposition of the mean score that we developed to provide a compact and informative description of the key features of the action distribution characterizing individual strategies and directly related to performance.

3.2 Validity of the Hessian-based decomposition

The Hessian-based performance analysis that we propose 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 do not contribute to the action score approximation in (26). To validate this assumption, we tested whether the score-relevant variability was not too large compared to the local tolerance and whether (27) can be considered as an acceptable model of the mean action score. Fig 6 shows the relationship described by (27) between the (sample) mean squared error 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). 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.

3.3 Score-relevant variability identified by Hessian-based decomposition of mean score

The Hessian-based decomposition allows to characterize the structure of individual variability and its relation with the local geometry of the score, overcoming the limitations of a qualitative description of the action distribution and the score function, which is challenging when the score is defined over a high-dimensional space of action variables. Similar to the UCM and GEM method, which use Jacobian matrices to split motor variability along task-relevant and task-irrelevant directions, we use the Hessian matrix to quantify *score-relevant* variability affecting the mean score. In this section, we illustrate how, across participants, a few eigenvectors of the Hessian matrix (principal sensitivity directions) associated with the largest eigenvalues identify the score-relevant directions that determine whether action variability affects the mean score or not. In the following sections we will then focus on how a few terms of the Hessian-based decomposition of the mean score (α, β, θ, τ, and η) characterize the different strategies employed by each participant, i.e. individual action distributions with specific mean and covariance. Before, exploring how individual strategies can be characterized through the terms of the Hessian-based decomposition, in the following we characterize the structure of the Hessian and the action covariance matrices, which in turns define the Hessian-based terms.

Fig 7A shows the distributions, across participants and for each target, of the eigenvalues or singular values of the Hessian matrix. Because in our throwing task the outcome space is two-dimensional, the solution manifold is a four-dimensional manifold embedded in the six-dimensional action space (see S1 Appendix C). Hence, on the solution manifold, the Hessian matrix will only have two non-zero singular values, whose associated eigenvectors or singular vectors defines a score-relevant plane. The corresponding eigenvalues quantify the sensitivity of the score function along that direction. Away from the solution manifold however, the Hessian matrix is also influenced by the non-linearities introduced by the mapping between actions and outcomes. Hence, for participants that do not have optimal mean actions, the Hessian matrix can have additional singular values which are different from zero. Fig 7A shows, however, that the contribution of the third and fourth singular value is negligible compared to the first two. Finally, focusing on the first two singular values, we notice that the sensitivity is slightly anisotropic, with the first sensitivity, on average, about 10% higher than the second. Furthermore, the first sensitivity shows the largest variability across participants and target conditions.

The two principal sensitivity directions $\boldsymbol{u}_1^H$, $\boldsymbol{u}_2^H$ define locally, i.e. around the mean action, a *sensitivity plane* embedded in the six-dimensional action space of the release parameters. Fig 7B shows the distributions of the principal sensitivity directions across participants and for each target. Across targets, the first principal sensitivity direction $\boldsymbol{u}_1^H$ is dominated by the vertical components (both position and velocity) of the release parameters and the second principal sensitivity direction is instead dominated by the lateral components (x_0, v_{x0}) of the release parameters.

In terms of action variability, for all participants and for all target conditions three principal components were able to explain 95% of the total variation, as shown in Fig 7C. Across participants and target conditions, the eigenvalue of the first principal component shows large variability across participants. It should be noted that the covariance matrix and, thus, the number of principal components is coordinate-dependent and in our scenario the action vectors contain both position and velocity variables which have different units. To assess the robustness of the estimation of the dimensionality of the action variability, we performed a principal component analysis on the correlation matrix rather than on the covariance matrix. The analysis of the eigenvalues of the correlation matrix confirmed that across participants and target conditions there were no more than three eigenvalues grater than one [15], supporting the conclusion that three components are sufficient to adequately describe the action variability.

Fig 7D illustrates the distribution of the first three principal variability directions $\boldsymbol{u}_1^\Sigma$, $\boldsymbol{u}_2^\Sigma$, $\boldsymbol{u}_3^\Sigma$, i.e. the eigenvectors of the covariance matrix associated to the three largest eigenvalues, indicating the directions along which most of the variation occurs. As for the principal sensitivity directions, a number of features of the principal variability directions are consistent across targets. For all four targets, the first direction captures a negative correlation between longitudinal (v_y) and vertical (v_z) release velocities and the second direction is dominated by the lateral velocity (v_x) but with variable contribution of the other velocity components. The distribution of all three directions, in larger measure for the second and third direction, however, are broader than those for the principal sensitivity directions, indicating that there are larger inter-individual differences in the structure of the action variability than in the sensitivity of the score around the mean action.

The distributions of principal sensitivity directions and principal variability directions illustrated in Fig 7B and D characterize features of the throwing strategy consistent across targets and participants. However, it is the selection of specific directions and, even more, their geometric relationship that determines the performance of individual participants. For example, two strategies with identical mean action, which implies identical sensitivity, may have different mean scores because of different alignments of the principal variability directions with respect to the principal sensitivity directions. Fig 7E shows the absolute values of the scalar products between the 6 pairs of principal sensitivity directions ($\boldsymbol{u}_1^H$, $\boldsymbol{u}_2^H$) and principal variability directions ($\boldsymbol{u}_1^\Sigma$, $\boldsymbol{u}_2^\Sigma$, $\boldsymbol{u}_3^\Sigma$) in all 20 participants. Such complex pattern of scalar products between principal sensitivity and variability directions highlights a remarkable inter-individual variability in the relationship between action variability and score sensitivity. In the next section we show how a few parameters deriving from the Hessian-based decomposition of the mean score compactly describe all performance related features of the individual action distribution and allow to fully characterize how performance depends on such features.

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

Fig 8 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 15) across participants and relative to target T1. To visualize the inter-individual differences in terms of these five parameters, we split the parameters into three pairs: the α-β plane in Fig 8A, the η-β plane in Fig 8B and the τ-θ plane Fig in 8C.

Moreover, since β is the product of θ/τ and η and $log(\beta) = log(\theta/\tau) + log(\eta)$, we also consider the $log(\eta)$-$log(\theta/\tau)$ plane in Fig 8D to visualize the additive contribution of these terms to the *tolerance-variability* index.

The $\alpha - \beta$ plane and the trade-off between bias and tolerance-variability

The $\alpha - \beta$ plane (Fig 8A) shows the performance of each individual thrower and its decomposition in terms of α and β. In this plane, Eq (27) defines a family iso-performance (same mean score) lines. The origin of the plane ($\alpha = 0, \beta = 0$) defines the optimal strategy, i.e. the one which achieves on average zero penalties. The closer a participant is to this point, the better its performance, as for $P18$ and $P4$. Participants with $\alpha \approx 0$ such as $P18, P2, P1$ have their mean action on (or very close to) the solution manifold, i.e. the set of actions which results into zero penalties. For these participants, deviations from the optimal performance $E[\pi] = 0$ are only due to the *tolerance-variability* index β and hence to the geometric relationship between action variability and score tolerance. Compared to $P18$, $P1$ has a higher tolerance-variability index (β) and hence on average collects more penalties. Conversely, while $P5$ and $P15$ have the same β, hence the same contribution of tolerance-variability, $P15$ perform worse than $P5$ due to a higher α. In other words, the mean action of $P15$ is inaccurate (or not optimal), i.e. it does not belong to the solution manifold, and hence it increases, by α, its mean score $E[\pi]$. Lastly, notice that participants such as $P1, P15$ and $P11$ are examples of *iso-performing* (same mean score) throwers that trade off smaller β for higher α and vice versa. $P1$ and $P11$ have $\alpha \approx 0$, their mean release action is optimal and does not influences the mean score. Conversely, $P15$, has a smaller β, hence the tolerance-variability score has less influence on the mean score, however this is compensated by a larger α, hence a less optimal mean action and a large systematic error error between target position and average outcome.

In sum, the decomposition of performance as $\alpha + \beta$ allows to characterize individual throwing strategies in the $\alpha - \beta$ plane in terms of trade-off between bias (α) and tolerance-variability (β).

Noise, tolerance and alignment

Fig 8B shows the distribution across participants of the total variation, or uncorrelated noise η, and its relation with the tolerance-variability index β. According to (13), participants lying on a given straight line, passing through the origin ($\eta = 0, \beta = 0$), have the same alignment-to-tolerance ratio $\frac{\theta}{\tau}$, which is equal to the slope, and their β index depends only on the amount of stochastic action variability η. Lower slopes correspond to more efficient strategies, as the slope is minimized by high τ values, i.e. when the mean action is in a high tolerance region of the action-to-score mapping, and by low θ values, i.e. by action distributions less aligned with the more sensitive directions.

There are many differences across participants. The best throwers such as $P8, P4, P13$ have relatively little action variability η, relatively little alignment-to-tolerance ratio and hence relatively little β score. Conversely, less skilled throwers such as $P20, P3$, and $P19$ lie on high-slope lines, populating regions with an intermediate level of noise but relatively high β. These participants have a relatively high alignment-to-tolerance ratio meaning that they are releasing actions from a very sensitive region (low tolerance) of the score and/or that principal variability

directions and principal sensitivities are aligned as shown in the τ-θ plane (Fig 8C). Intermediate and iso-performing throwers can be found all over the β-η plane. Notice that despite $P11$ is about three times more variable than $P1$, they have similar β. This is because $P11$ releases actions from a very tolerant region and because it directs stochastic perturbations along directions which do not affect the score (small θ). Lastly notice that in the τ-θ plane participants with higher tolerance τ tend to also have lower alignment θ. Such negative correlation indicates that participants can more easily orient their action covariance along less score-sensitive direction of the action space in regions where the tolerance is higher. However, such relationship between τ and θ is not fixed, as participants with the same or similar τ show a range of θ values (e.g. $P10$, $P16$, and $P3$ with low τ and $P19$, $P2$, and $P5$ with intermediate τ). Moreover, the range of values for τ and θ (hence the scale of the axes in Fig 8C) is different, with τ ranging from 0.32 to 0.42 and θ ranging from 0.02 to 0.13, indicating that alignment contributes more than tolerance in determining β, and thus performance, for a given amount of noise.

The $log(\eta) - log(\theta/\tau)$ plane and the trade-off between noise and alignment-tolerance

We can compare the relative contribution of alignment and tolerance with respect to noise considering the $log(\theta/\tau)$-$log(\eta)$ plane (Fig 8D), as in the log-log plot the two factors contribute additively to $log(\beta)$. In this plane, participants with the same β lie on the same line with slope -1 and participants with better tolerance-variability index, and hence better performance for a given α, lie on lines closer to the origin. This plane then reveals a second trade-off that characterizes iso-performance strategies in addition to the trade-off between α and β. Thus, this plane highlights how the individual action distribution (i.e. action mean and covariance) affects β in addition to the direct contribution of α (score of mean action). The choice of mean action determines the geometry of the task (i.e. the local tolerance τ and the most score sensitive directions of the action-to-score mapping) but the two key features of the action distribution affecting β are the overall amount of noise (η) and the alignment of the action covariance with the most score sensitive directions (θ).

3.5 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 9 shows the distributions across participants and targets of the parameters of the Hessian-based decomposition in the same combinations of parameters as in Fig 8, 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 α-β plane (Fig 9A), a few participants show large across-target variations (e.g. $P14$, who has 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 9B) and $log(\eta)$-$log(\theta/\tau)$ (Fig 9D) planes. Finally, in the τ-θ plane (Fig 9C) about half of the participants occupy the region with intermediate values of τ.

A 2-way ANOVA with participant and target as factors (Table 2) on 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 τ. Thus, τ is the only parameter that is not consistent across targets. This is not surprising since τ 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.

parameter	P-value (participant)	P-value (target)
α	$***$	0.07
β	$***$	0.39
θ	$***$	0.20
τ	$***$	$***$
η	$***$	0.77

Table 2. Robusteness of decompisition parameters across targets. P values of 2-way ANOVA with participant and target as factors. ***: $p < 10^{-3}$.

3.6 Dependence on throwing speed

As some of the geometric features of the action-to-score function captured by τ and affecting θ appear to depend on the magnitude of the release velocity (see Fig 5), whose largest component is longitudinal (v_{y0}), we investigated the dependence of the mean score and of the parameters of the Hessian-based decomposition as a function of release speed (i.e. $||\boldsymbol{v_0}||$). Furthermore, we wondered if speed is consistent across targets and at what extent it determines the individual strategy.

Fig 10A shows the distribution of the mean release speed for each target and participant. For each participant, sorted according to overall mean speed, each circle indicates the mean speed for one target and the box highlights the individual performance (with the same color coding as in Fig 9) and the range of mean speeds (with the horizontal line within each box representing the overall mean). Mean speed is broadly distributed across participants (range of overall mean: 8.5-13.3 m/s) but highly consistent across targets (range of across-target standard deviation: 0.01-0.39 m/s). The hottest colors, concentrated in middle of the plot, highlight how participants throwing with intermediate speeds (e.g., $P1$, $P4$, $P8$, $P12$, and $P18$) have on average better performance.

To better assess the effect of speed on performance, we averaged the mean score and its Hessian-based decomposition parameters in the three speed intervals indicated by the gray-shaded areas in Fig 10A). Fig 10B shows the average values of the mean score ($\alpha + \beta$), the tolerance-variability index (β), the ratio between alignment and tolerance (θ/τ), and noise (η), in the three speed intervals. Interestingly, not only the mean score, but also β shows a U-shaped

trend with speed. This implies that participants throwing with intermediate mean speeds tend to perform better and are more efficient in keeping the tolerance-variability index low.

Such U-shaped speed-dependence of mean score and β results from the different dependence on speed of alignment, tolerance, and noise, as it can be seen in Fig 10D-F, which show the dependence across individual participants and targets. The contribution to the β index of alignment and tolerance, θ/τ, tends in fact to decrease with speed, while the contribution of noise η have the opposite trend, increasing significantly at highest speeds.

Despite these overall trends, however, the distributions of θ and η exhibit a large variability, particularly in the intermediate range of speeds. This implies that the average speed of a thrower cannot fully predict the individual throwing strategy (i.e. the combination of Hessian-based parameters), neither how β contributes to the individual overall performance. For example, participants with similar intermediate mean speeds (e.g. $P6$, $P13$, and $P20$) may have dissimilar β, mostly due to differences in θ and η, as the range of variation of τ at equal speed is smaller. On the other hand, participants with similar mean speed, can achieve similar values of β with relative amounts of θ and η that deviate differently from the average speed dependence. This is the case for the two best performing subjects ($P4$ and $P18$), with $P18$ compensating his higher action variability with a better alignment.

3.7 A low-dimensional representation of individual strategies by projection onto the principal sensitivity plane

Finally, we introduce a low-dimensional representation of individual action distributions, based on the same observation underlying the Hessian-based decomposition, i.e. that the directions of maximal score sensitivity are described by the first few eigenvectors of the Hessian of the action-to-score function and performance depends on how the directions of maximal action variability are aligned with the directions of maximal sensitivity. Such representation allows to visualize in a single two-dimensional plot all the key features characterizing individual throwing strategies and described by the indicators derived from the Hessian-based decomposition.

The bar plots in Fig 11A shows the individual principal sensitivity directions and principal variability directions for the same five representative participants illustrated in Fig 5. Notice that, as also shown in Fig 7, there is little difference in terms of principal sensitivity directions (*green bars*) across participants, while individual differences can be appreciated in terms of action variability, in particular the directions along which variability is the highest across participants (*blue bars*). Due to high dimensionality of the problem, however, it is difficult to understand from these bar plots how each principal component contributes to the score-relevant variability and hence to the β index. However, because in our scenario the Hessian is locally dominated by two principal sensitivity directions, we can plot on a plane both the score relevant dimensions and the score relevant variability, and hence visualize the contribution of each principal component of a strategy.

Fig 11B shows the principal sensitivity plane in a neighborhood of the mean release action (*blue large circle*) of each participant and the distribution of actions as score-relevant perturbations of the mean release action (*red small circles*). To represent the sensitivity plane and action score in Fig 11B, we plot the first and second principal sensitivity directions $\boldsymbol{u}_i^H$ along the horizontal and the vertical axes, respectively. These two 6-dimensional axes are centered on

the mean action of each participant. Each point $\boldsymbol{p} = \left[\ p_1\boldsymbol{u}_1^H, p_2\boldsymbol{u}_2^H\ \right]$ on the plane represents a perturbation of the average action $\bar{\boldsymbol{a}}$ tangential to the sensitivity plane. In this way it is possible to visualize the best (local) planar representation of the score (*gray-shaded areas*). Notice that the white disk is smaller for less tolerant participants such as $P10$ and $P1$ compared to more tolerant participants such as $P18$ and $P11$. Also notice that the principal sensitivity frame (*green lines*) for $P15$ has a large shift compared to the center of disk, highlighting once again that mean action of this participant is not optimal and therefore does not belong to the solution manifold (the center of the disk in this representation).

Fig 11C shows how the alignment of the individual principal variability directions relative to the principal sensitivity plane contributes to the score-relevant variability and how the local sensitivity amplifies its detrimental effect for performance. The *magenta thin lines* are the principal variability directions (eigenvectors of the action covariance matrix) *projected* on the principal sensitivity plane and therefore their length provide a visual representation of the alignment between principal sensitivity and principal variability direction. By rescaling each projected eigenvector by their respective eigenvalues, the *red thick lines* show the contribution of each principal component of variability to the score relevant variability. Notice that, on the principal sensitivity plane, $P10$, which is one the most aligned participant, has the largest projection of the principal variability directions (i.e. the longest magenta thin lines), compared to less aligned participants such as $P18$ and $P11$. However, when scaled by the amount of variability in each direction (thick red lines), the projected principal variability directions of $P10$ have a similar magnitude as those of $P1$ and $P11$.

We can then represent the score-relevant variability (Eq 19) and its modulation due to local sensitivity (B matrix in Eq 20) as ellipses. The *yellow ellipses*, representing the one-standard deviation covariance matrix projected onto the principal sensitivity plane ($\Sigma_H = U_2^{H^T}\Sigma U_2^H$), show the score-relevant variability of each participant. The *gray ellipses*, representing one-standard deviation covariance matrix projected onto the principal sensitivity plane and scaled by the eigenvalues of the Hessian ($B = \Lambda_2^{H^{\frac{1}{2}}}\Sigma_H\Lambda_2^{H^{\frac{1}{2}}}$), show the effect of tolerance, and hence how local sensitivity amplifies score-relevant variability. The gray level of those ellipses represents the tolerance-variability index β, which is directly related to their size ($\beta = \frac{1}{2}trace(B)$). Notice that $P10$, $P1$ and $P11$ have a similar score-relevant variability (area of the yellow ellipses), however, $P10$ has a larger (and darker) gray ellipse due to a larger sensitivity of the score.

In sum, the projection of the action distribution and the principal variability directions onto the principal sensitivity plane provides a low-dimensional yet informative representation of Hessian-based decomposition indicators. The bias α is displayed by the offset of the solution manifold. The tolerance-variability index β corresponds to the size and shading of the gray ellipses. The tolerance τ is evident from the difference in size between the yellow and gray ellipses. The alignment θ can be inferred from the length of the thin magenta lines. Finally, the noise η determines the relative length of the red thick lines with respect to the thin magenta lines.

4 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 $\bar{\boldsymbol{a}}$ and action covariance Σ^a across multiple trials. We showed that performance, defined as the mean score, can be approximated as the sum of two components: the score of the mean action ($\alpha(\bar{\boldsymbol{a}})$) and a tolerance-variability index ($\beta(\bar{\boldsymbol{a}}, \Sigma^a)$). 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 the actions variability (stochastic noise) and by the geometry of the action-to-score function (determining the sensitivity to noise as a result of the non-linearities around the mean action and their alignment with the directions of largest variability). Such index results from the product of three terms: (i) the *total action variability* (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 action-to-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 (quantified in terms of the trace of the Hessian); (iii) the *alignment* (θ), a scalar measure of the relative orientation between the directions of curvature of the action-to-score mapping (indicated by the eigenvectors associated to the non-zero eigenvalues of the Hessian matrix computed at the mean action) and the directions of maximum variability of the actions (indicated by the eigenvectors associated to the largest eigenvalues of the action covariance matrix). Thus, these five parameters 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 $\alpha, \beta, \eta, \tau$ and θ parameters (Fig 8) and most of these differences were consistent across targets (Fig 9). In line with previous works focusing on low-dimensional throwing tasks [21, 5], 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. 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). Finally, participants consistently threw with different speeds, and speed affected the parameters of the

Hessian-based decomposition, but did not fully predict the individual throwing strategy. The low performance and the small variability of the decomposition parameters observed at low and high speeds may be due to the intrinsic constraints of the throwing task geometry (high θ/τ at low speed) and in the signal-dependent nature of the noise (high η at high speed). In contrast, at intermediate throwing speeds it may be possible to better trade off the different parameters for optimizing performance.

Such interplay between variability, speed, and geometric features of the action-to-score mapping can be observed in several pairs of action variables (e.g. in the v_y-v_z plane, Fig 5) or, in a unique and informative way, in the projection of the action vectors on the subspace spanned by the first two eigenvectors of the Hessian matrix (Fig 11). Thus, our decomposition provides a compact yet detailed characterization and visualization of the features of individual action distributions affecting performance.

4.1 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 [24, 17, 21, 13], our method highlights the key role of the geometry of the mapping between actions and outcomes or scores to assess how action variability affects performance. Differently from first-order methods such as UCM, GEM, or the more recent approach in [32], 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 (5) and (6), 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 method [21, 4, 28, 27]. The method has 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. 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 for generating surrogate data, 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 detailed in the Appendix (see S1 Appendix A), 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. Importantly, all the three terms of the TNC decomposition depend on both the Hessian and the action covariance. As an advantage, however, the TNC method does not rely on any assumption on the action-to-score mapping, such as smoothness and adequateness of a second-order approximation.

A fundamental difference between GEM, and 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. As summarized in Table 3, central to the GEM analysis is a Goal-function $e(\boldsymbol{a}) = \boldsymbol{x}_T - \boldsymbol{f}(\boldsymbol{a})$ that expresses the error between a desired goal/target $\boldsymbol{x}_T$ and the outcome $f(\boldsymbol{a})$ of a given action a. By linearizing this error/goal function around the mean action of a strategy, GEM quantifies the contribution of *tolerance*, *noise* and *alignment* to the overall task/goal-level variability, i.e $var(\boldsymbol{e})$. 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 (27).

	GEM	TNC	Hessian-based
performance variable	$e(a, x_T) = x_t - f(a)$	$\pi = s(e(a, x_T))$	$\pi = s(e(a, x_T))$
performance statistics	$var(e)$	$E[\pi]$	$E[\pi]$
decomposition approach	linearization of e	surrogate datasets	quadratic approximation of s

Table 3. Comparison between different methods relating performance, variability and task geometry. In GEM analysis, *tolerance*,*noise* and *alignment* characterize the effect that action variability and error sensitivity have on goal-level error variability. Differently, the TNC and the Hessian-based decomposition target the *mean* of a cost function of the error, i.e *the expected score/loss*, as shown in Fig.2.

4.2 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 tasks involving as action outcome a spatial position with respect to a goal, such as the arrival position of a projectile on a target board with respect to the center of the target, and requiring to minimize the distance from the goal, 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 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 Σ^a and the Hessian H^a 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 [2]). 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 S1 Appendix B). 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] and of UCM and GEM methods [29]. While such dependence may provide an opportunity to evaluate the role of different coordinate systems for

control [20], it has also been noticed that geometric properties of the action-to-score function such as the solution manifold do not depend on coordinates [29]. 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 [14]. Furthermore, the variability analysis is shown to be coordinate-independent as the characterization is performed in the eigenspace of the error-correcting controller matrix.

4.3 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 [25]; and a model-free system in which learning is driven by reinforcement and punishment of successful/erroneous actions [9, 3]. 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 this two learning mechanisms, remains poorly understood.

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 $\boldsymbol{f}$, 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.

4.4 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. Finally, participants consistently differed in throwing speed and speed affected θ/τ and η in opposite directions, because of the action-to-score geometry and signal-dependent noise, such that the best performances were observed at intermediate speeds. However, speed did not fully determine those parameters. 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.

5 Appendix

A. Relation to TNC approaches

A.1 Relation to Müller & Sternad 2004

In this section we show that, when the score function is smooth, and the action distribution is sufficiently localized, it is possible to derive (Hessian-based) analytic expressions to isolate the three components of the TNC approach by Müller and Sternad [21], shown in Fig.12. Given two experimental strategies, such as S_A and S_B in the figure, the TNC method requires the generation of surrogate data-sets, S_A^0, S_B^0 and S_A^{sh}, to decompose the difference in expected score $\Delta\bar{\pi} = \bar{\pi}(S_B) - \bar{\pi}(S_A)$, into the sum of four independent components: ΔC_1, or *covariation*, is the difference in expected score between the strategy S_A and the (surrogate) strategy S_A^0, that is obtained by removing (via random permutations) any linear/non-linear correlation between the variables of the data-set S_A. Similarly, for S_B, a surrogate uncorrelated data-set S_B^0 is used to quantify the difference in performance ΔC_2 due to covariations in S_B. Notice that S_A^0 and S_B^0 have the same mean ($\bar{\boldsymbol{a}}$ and $\bar{\boldsymbol{b}}$, respectively) as their original data-sets and only differ with respect to their original data-sets in terms of variability. A third surrogate data-set S_A^{sh} is

generated by *shifting* the location of S_A^0 (i.e. $\bar{\boldsymbol{a}}$) to the average location $\bar{\boldsymbol{b}}$ of the S_B data-set. The *tolerance* component is hence quantified as $\Delta T = \bar{\pi}(S_A^{sh}) - \bar{\pi}(S_A^0)$, has the two data-sets have same variability and differ only in terms of their average location. Lastly, the *noise* component is extracted as the difference in average performance between the surrogate data-set S_B^0 and S_A^{sh}, i.e. $\Delta N = \bar{\pi}(S_B^0) - \bar{\pi}(S_A^{sh})$.

Assuming that motor strategies are drawn from a localized distributions, $S_A = \{\bar{\boldsymbol{a}}; \Sigma_A\}$ and $S_B = \{\bar{\boldsymbol{b}}; \Sigma_B\}$, our method allows the estimation of all four components without using surrogate data-sets and random permutation. In this case, the covariance matrices Σ_{A_0} and Σ_{B_0} of the uncorrelated strategies S_A^0 and S_B^0, can be simply computed as diag(Σ_A) and diag(Σ_B), i.e. as the matrix of the diagonal elements (variances) of Σ_A and Σ_B, respectively. Knowing the Hessian matrix $H_{\bar{\boldsymbol{a}}}$ and $H_{\bar{\boldsymbol{b}}}$ at the two locations $\bar{\boldsymbol{a}}$ and $\bar{\boldsymbol{b}}$, respectively, allows to approximate ΔC_1,ΔT, ΔN and ΔC_2, simply as:

$$\Delta C_1 = \bar{\pi}(S_A^0) - \bar{\pi}(S_A) \approx \left[\alpha(\bar{\boldsymbol{a}}) + \frac{1}{2} trace(H_{\bar{\boldsymbol{a}}}\Sigma_{A_0})\right] - \left[\alpha(\bar{\boldsymbol{a}}) + \frac{1}{2} trace(H_{\bar{\boldsymbol{a}}}\Sigma_A)\right] = \\ = \frac{1}{2} trace[H_{\bar{\boldsymbol{a}}}(\Sigma_{A_0} - \Sigma_A)]$$

$$\Delta T = \bar{\pi}(S_A^{sh}) - \bar{\pi}(S_A^0) \approx \left[\alpha(\bar{\boldsymbol{b}}) + \frac{1}{2} trace(H_{\bar{\boldsymbol{b}}}\Sigma_{A_0})\right] - \left[\alpha(\bar{\boldsymbol{a}}) + \frac{1}{2} trace((H_{\bar{\boldsymbol{a}}}\Sigma_{A_0})\right] = \\ = \alpha(\bar{\boldsymbol{b}}) - \alpha(\bar{\boldsymbol{a}}) + \frac{1}{2} trace[(H_{\bar{\boldsymbol{b}}} - H_{\bar{\boldsymbol{a}}})\Sigma_{A_0}]$$

$$\Delta N = \bar{\pi}(S_B^0) - \bar{\pi}(S_A^{sh}) \approx \left[\alpha(\bar{\boldsymbol{b}}) + \frac{1}{2} trace(H_{\bar{\boldsymbol{b}}}\Sigma_{B_0})\right] - \left[\alpha(\bar{\boldsymbol{b}}) + \frac{1}{2} trace(H_{\bar{\boldsymbol{b}}}\Sigma_{A_0})\right] = \\ = \frac{1}{2} trace[H_{\bar{\boldsymbol{b}}}(\Sigma_{B_0} - \Sigma_{A_0})]$$

$$\Delta C_2 = \bar{\pi}(S_B^{sh}) - \bar{\pi}(S_B) \approx \left[\alpha(\bar{\boldsymbol{b}}) + \frac{1}{2} trace(H_{\bar{\boldsymbol{b}}}\Sigma_{B_0})\right] - \left[\alpha(\bar{\boldsymbol{b}}) + \frac{1}{2} trace(H_{\bar{\boldsymbol{b}}}\Sigma_B)\right] = \\ = \frac{1}{2} trace[H_{\bar{\boldsymbol{b}}}(\Sigma_B - \Sigma_{B_0})]$$

Hence, it follows that the difference in expected performance between the two strategies can also be approximated as:

$$\Delta\bar{\pi} = \bar{\pi}(S_B) - \bar{\pi}(S_A) = \Delta C_1 + \Delta T + \Delta N + \Delta C_2 \approx \\ \approx \alpha(\bar{\boldsymbol{b}}) - \alpha(\bar{\boldsymbol{a}}) + \frac{1}{2} trace(H_{\bar{\boldsymbol{b}}}\Sigma_B) - \frac{1}{2} trace(H_{\bar{\boldsymbol{a}}}\Sigma_A) = \Delta\alpha + \Delta\beta \tag{29}$$

A.2 Relation to Cohen & Sternad 2009

An alternative TNC approach, the TNC-Cost, has been proposed by Cohen & Sternad [4] to overcome some of the limitations of the original TNC described in the Appendix A1. In the

novel approach, the parameters of the decomposition are no longer dependent on the sequence of calculation, and the components are quantified with respect to an optimized data sets derived by transformation of one of the features.

The T-Cost. The T-cost is the algebraic difference between the mean score of the experimental dataset and the mean score of an optimized dataset that differs from the first dataset only in terms of the mean action. More specifically, the action space is firstly discretized in a grid of n points that are used to shift the action distribution of the original dataset. Hence, for each grid point the original and the shifted dataset have the same action variability and only differs with respect to the mean action. For each grid points the method evaluates the mean score of each shifted dataset, such that the optimal one is the one which results in the minimal (i.e. best) mean score.

In the assumption that experimental dataset S_A has a small dispersion, we can apply our framework to calculate, analytically, the T-cost. In particular, let $\bar{\boldsymbol{a}}$ and Σ_A be the mean action and action covariance of the experimental dataset S_A, and let $\bar{\boldsymbol{a}}^*$ be the shifted mean of the optimized data set S_A^*, then the T-Cost can approximated as:

$$\begin{aligned}\text{T-cost} = \bar{\pi}(S_A) - \bar{\pi}(S_A^*) \approx \left[\alpha(\bar{\boldsymbol{a}}) + \frac{1}{2}trace(H_{\bar{\boldsymbol{a}}}\Sigma_A)\right] - \left[\alpha(\bar{\boldsymbol{a}}^*) + \frac{1}{2}trace(H_{\bar{\boldsymbol{a}}^*}\Sigma_A)\right] = \\ = \alpha(\bar{\boldsymbol{a}}) - \alpha(\bar{\boldsymbol{a}}^*) - \frac{1}{2}trace\left[(H_{\bar{\boldsymbol{a}}} - H_{\bar{\boldsymbol{a}}^*})\,\Sigma_A)\right] = \Delta\alpha + \Delta\beta.\end{aligned} \tag{30}$$

In the hypothesis that the original dataset had already an optimal mean action, i.e. on the solution manifold, and that the grid was sufficiently discretized such that also the optimal action belongs to the solution manifold, than the tolerance cost of Cohen & Sternad would correspond to a $\Delta\beta$ due to differences in the local geometry of the score between the experimental and the optimized datasets,i.e $H(\bar{\boldsymbol{a}}) \neq H(\bar{\boldsymbol{a}}^*)$.

Notice however, that the tolerance cost depends on the action covariance of the original dataset; hence it does not provide a unique measure of *tolerance/sensitivity* to errors of the score function. In other words, given two experimental datasets S_A and S_B with the same mean action $\bar{\boldsymbol{a}} = \bar{\boldsymbol{b}}$ but different covariances, one would measure a different Tolerance-cost for the two distributions, despite locally, the score has the same Hessians for both distributions,i.e. $H(\bar{\boldsymbol{a}}) = H(\bar{\boldsymbol{b}})$.

The N-Cost. The N-cost is the algebraic difference between the mean score of the experimental dataset and the mean score of an optimized dataset, or with an optimized *noise* compared to the original dataset. More specifically, the optimized dataset is obtained through a sequence of n steps which progressively shrunk the data points (the actions) of the original dataset towards its mean action. For each step, the mean score of the shrunk dataset is calculated and the optimal dataset is identified as the one with minimal mean score. This procedure however, does not preserve the mean action and therefore the optimized dataset does not only differs in terms of *noise* (action variability) but also has a shifted mean compared to the original dataset. In the assumption that experimental dataset S_A has a small dispersion, we can apply our framework

to calculate, analytically the N-cost:

$$\begin{aligned}\text{N-cost} = \bar{\pi}(S_A) - \bar{\pi}(S_A^*) \approx \left[\alpha(\bar{\boldsymbol{a}}) + \frac{1}{2}trace(H_{\bar{\boldsymbol{a}}}\Sigma_A)\right] - \left[\alpha(\bar{\boldsymbol{a}}^*) + \frac{1}{2}trace(H_{\bar{\boldsymbol{a}}^*}\Sigma_A^*)\right] = \\ = \alpha(\bar{\boldsymbol{a}}) - \alpha(\bar{\boldsymbol{a}}^*) + \frac{1}{2}\left[trace\left(H_{\bar{\boldsymbol{a}}}\Sigma_A\right) - trace\left(H_{\bar{\boldsymbol{a}}^*}\Sigma_{A^*}\right)\right]\end{aligned} \tag{31}$$

Therefore, the noise cost of Cohen & Sternad 2009, is not only affected by the variability in the original dataset Σ_A, but in general depends also on the local tolerance of the score as well as on the difference in score between the mean actions of the two datasets, i.e.$\alpha(\bar{\boldsymbol{a}}) - \alpha(\bar{\boldsymbol{a}}^*)$, given that, as also pointed out by author, the shrinking procedure can in general result in a mean action $\bar{\boldsymbol{a}}^*$ which is not optimal, or does not belong to the solution manifold $\alpha(\bar{\boldsymbol{a}}^*) \neq 0$. Furthermore, in the hypothesis that both datasets have optimal mean action, $\alpha(\bar{\boldsymbol{a}}) = \alpha(\bar{\boldsymbol{a}}^*) = 0$ and that the shrunk dataset has very little variability compared to the original one, the N-cost reduces to our β, i.e. $\frac{1}{2}trace(H_{\bar{\boldsymbol{a}}}\Sigma_A)$ highlighting the fact that this parameter does not in general provides a clear description of *noise* (meant as action variability) as it mixes all the contributions that are due not only to action distribution and variability, but also to the geometry of the score.

The C-Cost. The C-cost is the algebraic difference between the mean score of the experimental dataset and the mean score of an optimized dataset that differs from the first only in terms of correlation between action variables, hence having the same mean action and uncorrelated variability of the original dataset. Following our framework, in the hypothesis of localized distribution and smooth score function, the C-cost, or the difference between the original and the optimized dataset becomes:

$$\begin{aligned}\text{C-cost} = \bar{\pi}(S_A) - \bar{\pi}(S_A^*) \approx \left[\alpha(\bar{\boldsymbol{a}}) + \frac{1}{2}trace(H_{\bar{\boldsymbol{a}}}\Sigma_A)\right] - \left[\alpha(\bar{\boldsymbol{a}}^*) + \frac{1}{2}trace(H_{\bar{\boldsymbol{a}}^*}\Sigma_A)\right] = \\ = \frac{1}{2}trace\left[H_{\bar{\boldsymbol{a}}}\left(\Sigma_A - \Sigma_{A^*}\right)\right] = \frac{1}{2}trace\left(H_{\bar{\boldsymbol{a}}}T\right)\end{aligned} \tag{32}$$

where the T matrix has all zeros on the main diagonal given that the two datasets have the same uncorrelated noise.

B. Coordinate invariance

Approaches based on covariance matrices for the analysis of variability, such as the UCM, have often been criticized for their dependence on the choice of coordinates. Similar critiques have also been highlighted for the TNC approach that does not use (directly) covariance matrices: "for instance, one can always rotate the frame of reference to get variables that have zero covariance" [26]. Furthermore, it is well known that Principal Component Analysis is sensitive to co-ordinates, especially when the multivariate data contains variables with different units. For instance, in this work the action vector contains position that are measured in meters and velocities that are measured in ms^{-1}. Should we rescale the action space to have comparable variances between positions and velocities? Would scaling affect our results? Here we show that this is not the case and that both ours and the TNC approach [21] are invariant under affine coordinate transformations. In fact, scaling, rotations and translations, i.e. any affine transformation of the action space, does not only affect covariance matrices (and hence correlations

among variables) but also affects the performance manifold, in particular the structure of its Hessian.

Let's assume that we have two sets of coordinates $\{a\}$ and $\{b\}$ with which we can parameterize the n-dimensional action space $\mathcal{A} \subset \mathbb{R}^n$ and that the map $\boldsymbol{g}$ describe the relationship between the two coordinate system:

$$\boldsymbol{b} = \boldsymbol{g}(\boldsymbol{a}) \tag{33}$$

when the map is non-linear, its first-order approximation, around a point $\bar{\boldsymbol{a}}$ can be expressed as:

$$\boldsymbol{\delta}^b = J(\bar{\boldsymbol{a}})(\boldsymbol{\delta}^a) \tag{34}$$

where $J = \frac{\partial g}{\partial \boldsymbol{a}}$ is the $n \times n$ Jacobian matrix evaluated at $\bar{\boldsymbol{a}}$.

The score π is a scalar and therefore does not dependent on the choice of coordinates used to express the score function. In both coordinate systems we can write:

$$\pi = s^b(\boldsymbol{b}) = s^b(\boldsymbol{g}(\boldsymbol{a})) = s^a(\boldsymbol{a}) \tag{35}$$

By differentiating the last equality with respect to the $\{a\}$ co-ordinates, we find a well-known expression between the gradients in the two different co-ordinate systems:

$$\frac{\partial s^a}{\partial \boldsymbol{a}} = J^T \frac{\partial s^b}{\partial \boldsymbol{b}} \tag{36}$$

Differentiating again the above expression we can express the Hessian of the score in the two coordinate systems:

$$\frac{\partial^2 s^a}{\partial \boldsymbol{a}^2} = J^T \frac{\partial^2 s^b}{\partial \boldsymbol{b}^2} J + \left(\frac{\partial s^b}{\partial \boldsymbol{b}}\right)^T \frac{\partial J}{\partial \boldsymbol{a}} \tag{37}$$

Hence:

$$H^a = J^T H^b J + \frac{\partial J}{\partial \boldsymbol{a}}^T \boldsymbol{\nabla_b} s^b \tag{38}$$

When the average action belongs to the solution manifold ($\boldsymbol{\nabla_b} s^b = 0$), or when the change of co-ordinate is affine ($\frac{\partial J}{\partial \boldsymbol{a}}^T = 0$), the second term on the right hand-side is zero. In this case, the Hessian is a tensor and β does not depend on the co-ordinates used to parameterise the action space. In fact, let Σ^b, be the covariance of the action expressed in the $\{b\}$ coordinates, then, if the distribution is localized (small variability), the covariance in the $\{a\}$ co-ordinates can be estimated as:

$$\Sigma^a = J^{-1} \Sigma^b J^{-1^T} \tag{39}$$

and hence:

$$\beta = trace(\frac{1}{2} H^a \Sigma^a) = trace(\frac{1}{2} J^T H^b J J^{-1^T} \Sigma^b J) = trace(\frac{1}{2} H^b \Sigma^b) \tag{40}$$

where we have used the cyclic properties of the trace ($trace(ABC) = trace(CAB)$) to simplify the last equality.

Conversely, for highly non-linear change of coordinates, or for average actions that are 'far' from the solution manifold, the second term on the right-hand side of (38) may not be negligible. In such case, the Hessian looses its tensorial property, and β becomes a co-ordinate dependent measure:

$$\beta = trace(\frac{1}{2}H^a\Sigma^a) = trace(\frac{1}{2}H^b\Sigma^b) + trace\left(\frac{1}{2}\left(\boldsymbol{\nabla_b} s^b\right)^T \frac{\partial J}{\partial \boldsymbol{a}}\Sigma^b\right) \tag{41}$$

C. Non-zero Hessian's eigenvalues in the presence of redundant actions

In human motor control, the map between action and task variables represents a 'change of co-ordinates' $\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{a})$, which often is non-linear and redundant. This latter properties of the map, makes the Jacobian $J = \frac{\partial \boldsymbol{f}}{\partial \boldsymbol{a}}$ a rectangular matrix with n columns (dimension of the action space) and m rows (dimension of the task space). In this case, equation (38) becomes:

$$H^a = J^T H^x J + \frac{\partial J}{\partial \boldsymbol{a}}^T \boldsymbol{\nabla_x} s^x \tag{42}$$

where H^x and $\boldsymbol{\nabla_x} s^x$ are the $m \times m$ Hessian and the $m \times 1$ gradient of the task-score function, respectively, and H^a is the $n \times n$ Hessian of the action-score function. Again, either on the solution manifold, or for a linear map between actions and outcomes, the second term on the right-hand side disappears, and H^a will only have $m < n$ non-zero eigenvalues. Conversely, away from the solution manifold and for highly non-linear change of co-ordinates, the term $\left(\boldsymbol{\nabla_b} s^b\right)^T \frac{\partial J}{\partial \boldsymbol{a}}$ will in general affect the number of non-zero eigenvalues, as well as the symmetry and positive-definiteness of the Hessian matrix.

References

1. N A Bernstein. *On dexterity and its development.* 1996.

2. Domenico Campolo, Ferdinan Widjaja, Hong Xu, Wei Tech Ang, and Etienne Burdet. Analysis of Accuracy in Pointing with Redundant Hand-held Tools: A Geometric Approach to the Uncontrolled Manifold Method. *PLoS Computational Biology*, 2013.

3. Xiuli Chen, Peter Holland, and Joseph M. Galea. The effects of reward and punishment on motor skill learning. *Current Opinion in Behavioral Sciences*, 2018.

4. R. G. Cohen and D. Sternad. Variability in motor learning: Relocating, channeling and reducing noise. *Experimental Brain Research*, 2009.

5. Joseph P. Cusumano and Paola Cesari. Body-goal variability mapping in an aiming task. *Biological Cybernetics*, 2006.

6. Joseph P. Cusumano and Jonathan B. Dingwell. Movement variability near goal equivalent manifolds: Fluctuations, control, and model-based analysis. *Human Movement Science*, 2013.

7. Ashesh K. Dhawale, Maurice A. Smith, and Bence P. Ölveczky. The Role of Variability in Motor Learning. *Annual Review of Neuroscience*, 2017.

8. Carol A. Fowler and M.T. Turvey. 1 - skill acquisition: An event approach with special reference to searching for the optimum of a function of several variables. In George E. Stelmach, editor, *Information Processing in Motor Control and Learning*, pages 1 – 40. Academic Press, 1978.

9. Adrian M. Haith and John W. Krakauer. Model-based and model-free mechanisms of human motor learning. *Advances in Experimental Medicine and Biology*, 2013.

10. E. F. Harding, A. M. Matthai, and S. B. Provost. Quadratic Forms in Random Variables: Theory and Applications. *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 2006.

11. Christopher M. Harris and Daniel M. Wolpert. Signal-dependent noise determines motor planning. *Nature*, 1998.

12. Ya Ching Hung, T. R. Kaminski, Julie Fineman, Jane Monroe, and A. M. Gentile. Learning a multi-joint throwing task: A morphometric analysis of skill development. *Experimental Brain Research*, 2008.

13. Joby John and Joseph P. Cusumano. Inter-Trial Dynamics of Repeated Skilled Movements. pages 707–716, 2009.

14. Joby John, Jonathan B. Dingwell, and Joseph P. Cusumano. Error Correction and the Structure of Inter-Trial Fluctuations in a Redundant Movement Task. *PLoS Computational Biology*, 2016.

15. Henry F. Kaiser. The Application of Electronic Computers to Factor Analysis. *Educational and Psychological Measurement*, 1960.

16. K. P. Kording and D. M. Wolpert. The loss function of sensorimotor learning. *Proceedings of the National Academy of Sciences*, 2004.

17. Mark L Latash, John P Scholz, and Gregor Schöner. Motor Control Strategies Revealed in the Structure of Motor Variability. *Exerc. Sport Sci. Rev*, 30(1):26–31, 2002.

18. A. Maselli, A. Dhawan, M. Russo, B. Cesqui, F. Lacquaniti, and A. d'Avella. A whole-body characterization of individual strategies, gender differences and common styles in overarm throwing. *J Neurophysiol*, *submitted*, 2019.

19. Antonella Maselli, Aishwar Dhawan, Benedetta Cesqui, Marta Russo, Francesco Lacquaniti, and Andrea D'Avella. Where Are You Throwing the Ball? I Better Watch Your Body, Not Just Your Arm! *Frontiers in Human Neuroscience*, 2017.

20. Hermann Müller, Till D. Frank, and Dagmar Sternad. 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.

21. Hermann Müller and Dagmar Sternad. Decomposition of Variability in the Execution of Goal-Oriented Tasks: Three Components of Skill Improvement. *Journal of Experimental Psychology: Human Perception and Performance*, 2004.

22. K. M. Newell. Motor Skill Acquisition. *Annual Review of Psychology*, 1991.

23. Leslie C. Osborne, Stephen G. Lisberger, and William Bialek. A sensory source for motor variation. *Nature*, 2005.

24. John P. Scholz and Gregor Schöner. The uncontrolled manifold concept: Identifying control variables for a functional task. *Experimental Brain Research*, 1999.

25. R. Shadmehr, J. J. Orban de Xivry, M. Xu-Wilson, and T.-Y. Shih. Temporal Discounting of Reward and the Cost of Time in Motor Control. *Journal of Neuroscience*, 2010.

26. Jeroen B.J. Smeets and Stefan Louw. 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.

27. Dagmar Sternad. It's not (only) the mean that matters: variability, noise and exploration in skill learning. *Current Opinion in Behavioral Sciences*, 20:183–195, 2018.

28. Dagmar Sternad, Masaki O. Abe, Xiaogang Hu, and Hermann Müller. Neuromotor noise, error tolerance and velocity-dependent costs in skilled performance. *PLoS Computational Biology*, 2011.

29. Dagmar Sternad, Se Woong Park, Hermann Müller, and Neville Hogan. Coordinate dependence of variability analysis. *PLoS Computational Biology*, 2010.

30. Schmidt TD. and R A Lee. *Motor Control and Learning: A behavioural Emphasis*. 2005.

31. Elias B. Thorp, Konrad P. Kording, and Ferdinando A. Mussa-Ivaldi. Using noise to shape motor learning. *Journal of Neurophysiology*, 2016.

32. M. Venkadesan and L. Mahadevan. Optimal strategies for throwing accurately. *Royal Society Open Science*, 2017.

33. J-F Yang and JP Scholz. Learning a throwing task is associated with differential changes in the use of motor abundance. *Experimental brain research*, 163(2):137–158, 2005.

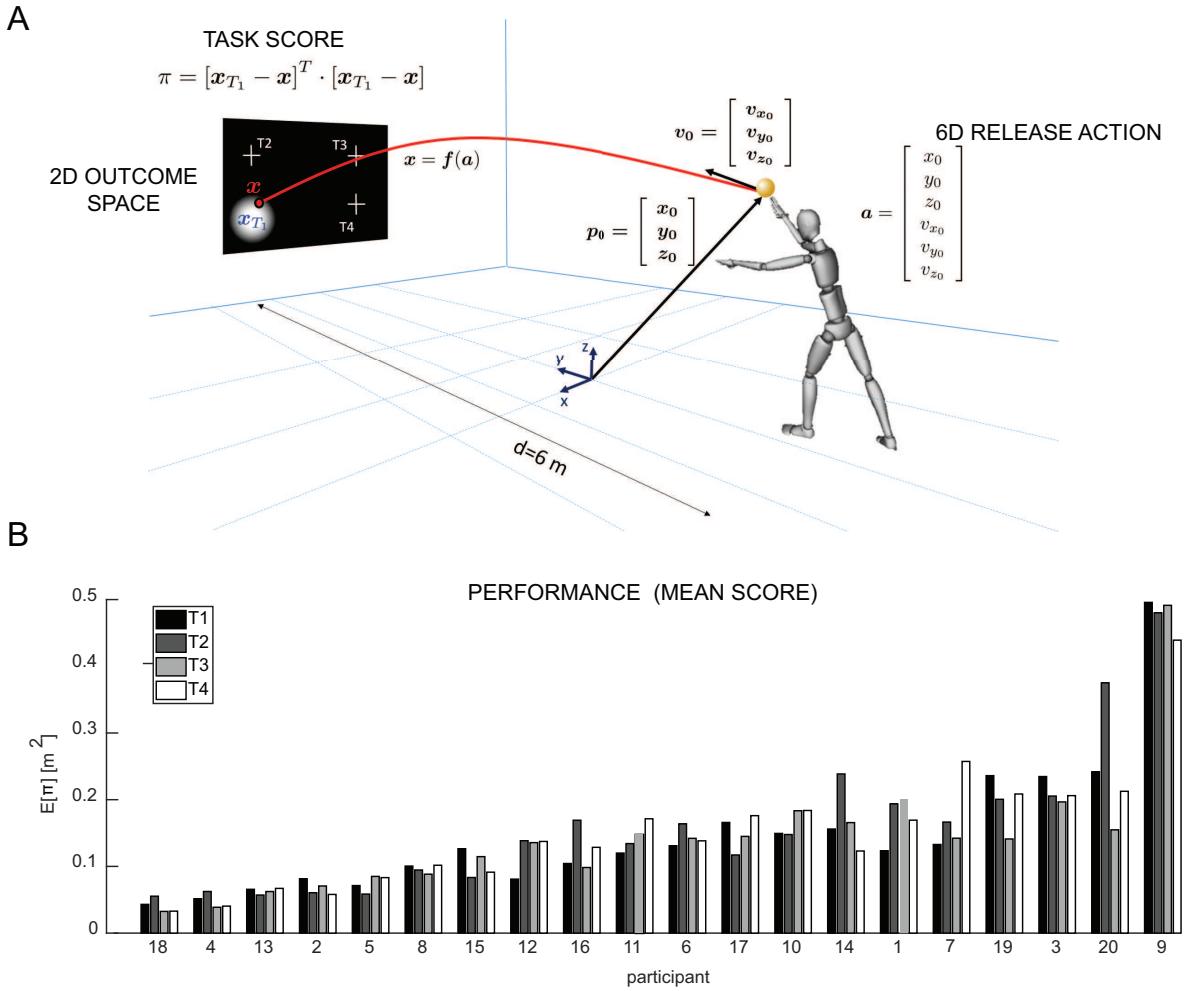

Figure 1. Throwing task and performance. (A) Schematic representation of our unconstrained overarm throwing scenario.(B) Performance or mean squared error across participants and targets.

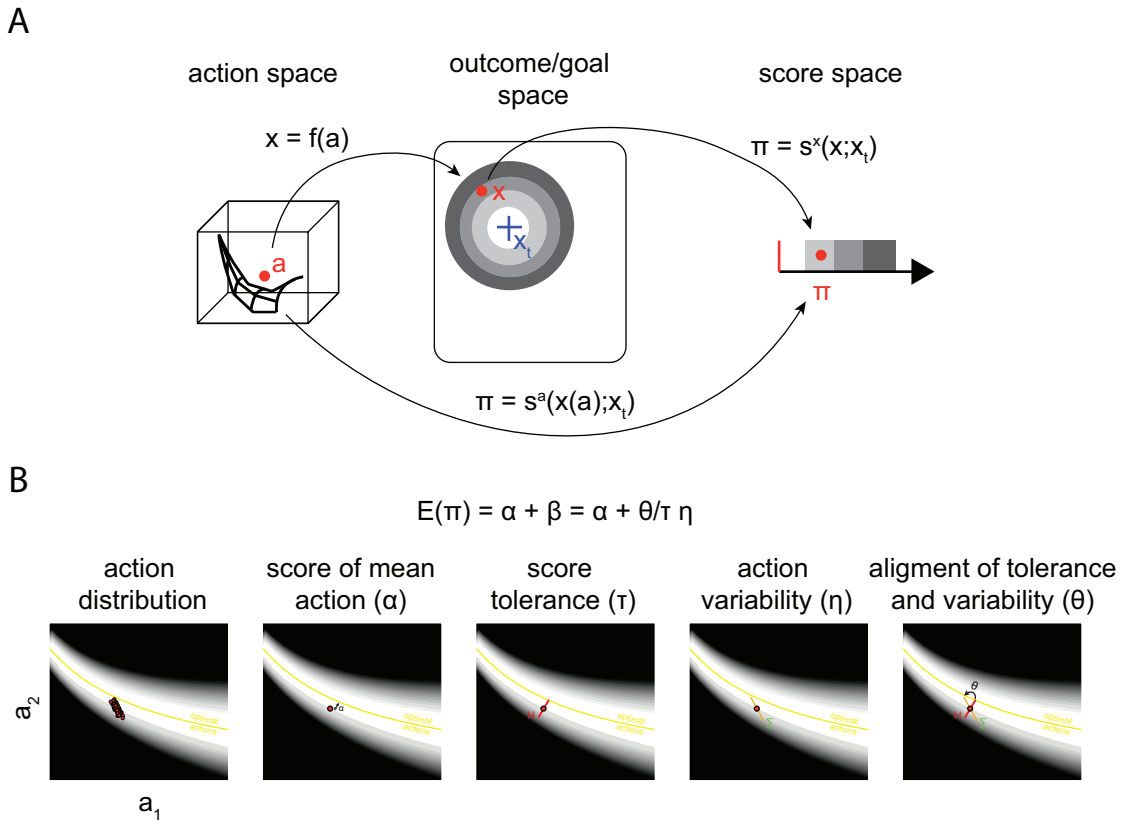

Figure 2. Definition of variables (action, outcome, and score) and schematic overview of the proposed approach to decompose the mean score (performance). (A) An action, represented by a vector $\boldsymbol{a}$ in a high-dimensional *action space* (illustrated by a *red* circular marker in a three-dimensional action space), leads to an outcome, $\boldsymbol{x} = f(\boldsymbol{a})$, i.e. a vector typically represented in a *outcome* or *goal space* with less dimensions. The outcome is then associated with a corresponding score, $\pi = s^x(\boldsymbol{x}, \boldsymbol{x}_T)$, 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, *left panel* has a distribution characterized by the mean (large red marker) and covariance (Σ, orange ellipse in the *fourth and fifth panels*). The tolerance of the score to variations of the actions around their means is characterized by the Hessian (H, red ellipse in the *third 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), the *uncorrelated noise* (η), and the *alignment* (θ).

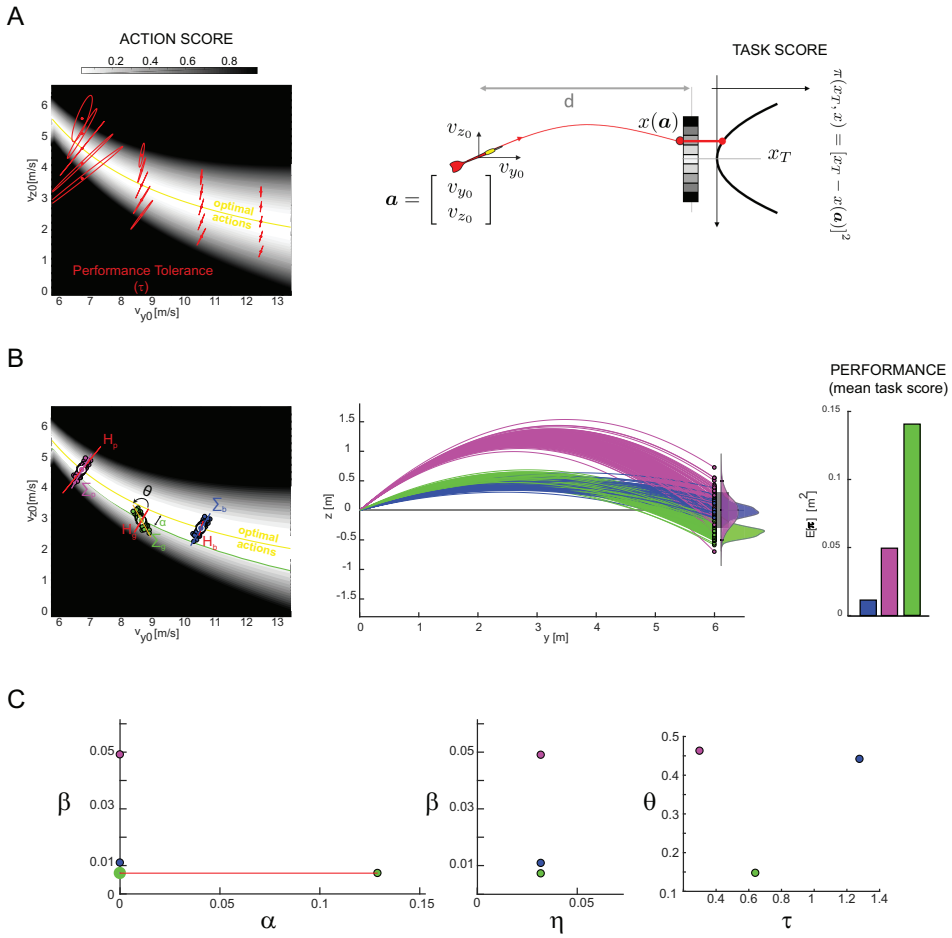

Figure 3. Simulated example of three different throwing strategies in a 2D throwing task. (A) Toy model of a throwing task involving a quadratic task-score. The action score and the solution manifold (yellow line) are shown in the left panel. The score sensitivity (the Hessian) is represented with red ellipses whose major axis represents, locally, the direction of maximum sensitivity (smaller noise tolerance). (B) Three (simulated) individual strategies (b=blue, g=green, p=purple), their task outcomes and accuracy performance (mean squared error) . (C) Hessian-based decomposition of the mean score of the three strategies according to (27). The blue and purple strategies have zero α (score of mean action) as the mean of the distribution of release parameters is on the solution manifold but they have different β; the green strategy has a non-zero α corresponding to a non-optimal mean action. Notice that the release strategies have been chosen to have the same level of noise η so as to highlight the effect of α, τ (tolerance) and θ (alignment) on the individual expected score. See the text for more details.

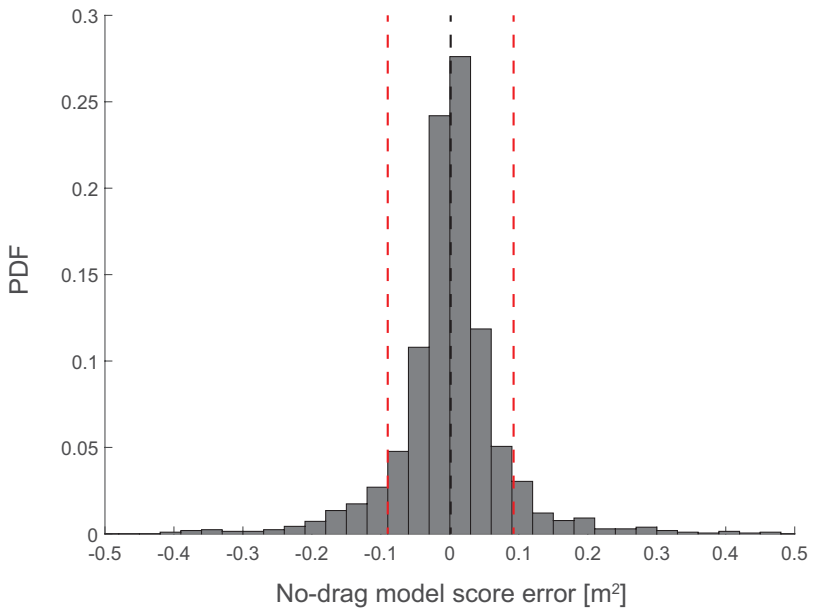

Figure 4. Error distribution between experimental performance and those predicted with the no-drag model in (26). Black and red lines indicate mean and $\pm 1SD$ of the distribution, respectively.

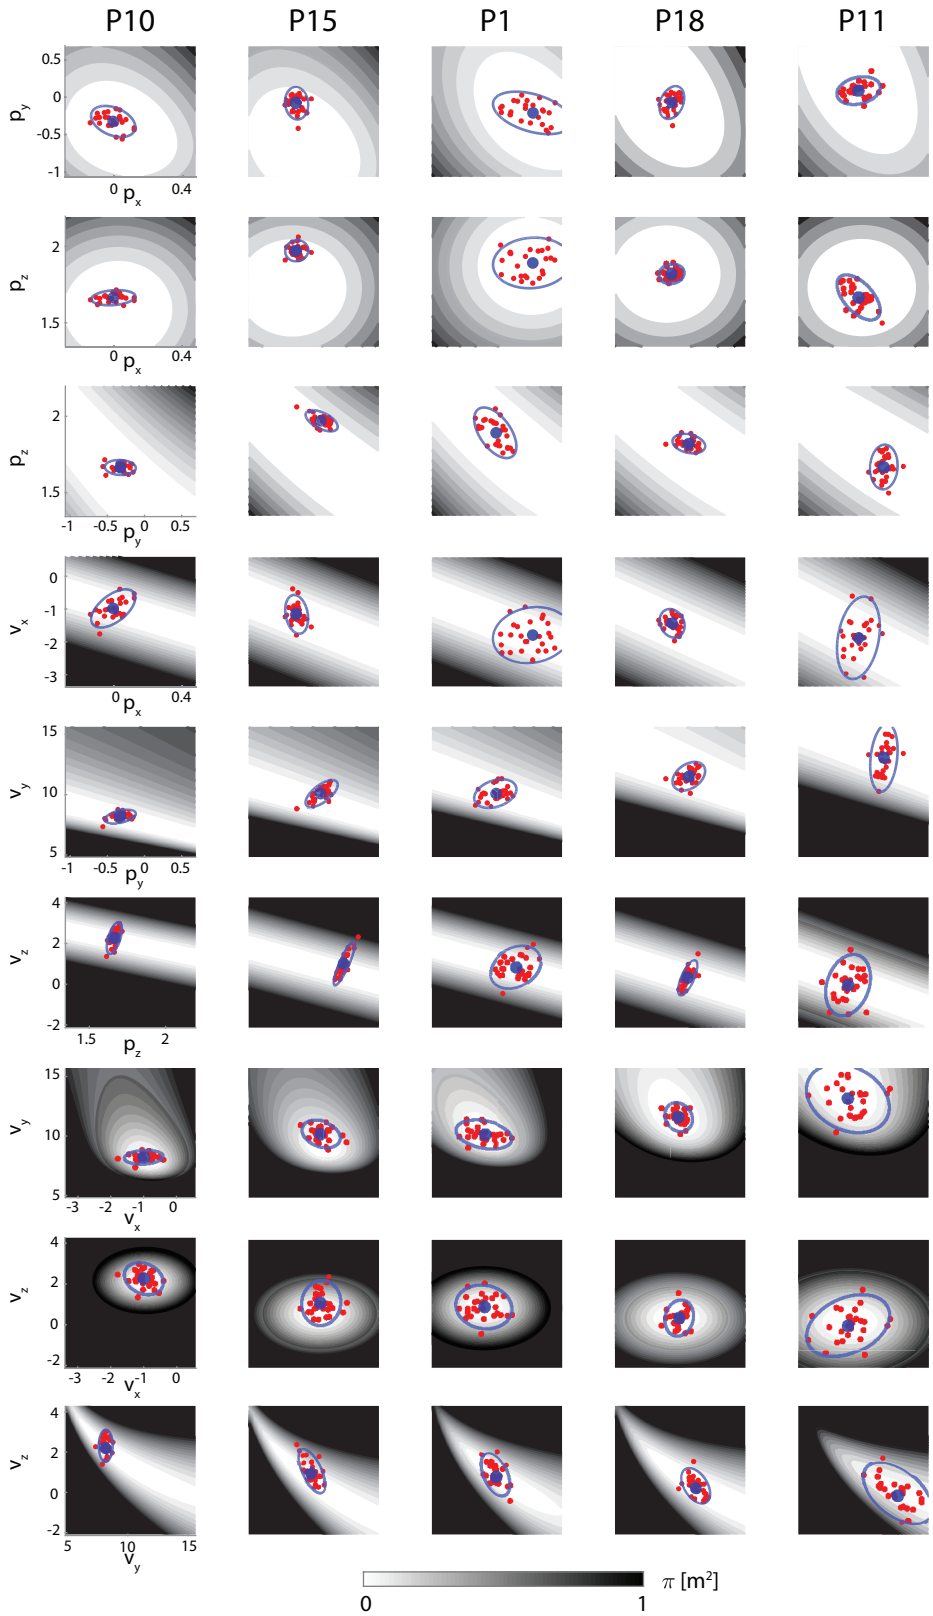
P10
P15
P1
P18
P11
p_y
p_x
p_z
v_x
v_y
v_z
π [m^2]
0
1

Figure 5. Examples of distribution of release position and velocity for five representative participants (target T1). Each row illustrates a pairs of release parameters, including all combinations of position variables (p_x-p_y, p_x-p_z, p_y-p_z), all combinations of position and velocity variables for each axis (p_x-v_x, p_y-v_y, p_z-v_z), and all combinations of velocity variables (v_x-v_y, v_x-v_z, v_y-v_z). Red circles represent release parameters of individual throws. Blue circles and ellipses represent mean and covariance (two standard deviations) of each parameter distribution. The gray level map shows the local score, as a function of the release parameters, underlying each throwing strategy. Note that differences in the background geometry of the maps reflects 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 stochastic perturbations. Participants have been sorted from left to right according to their average release speed.

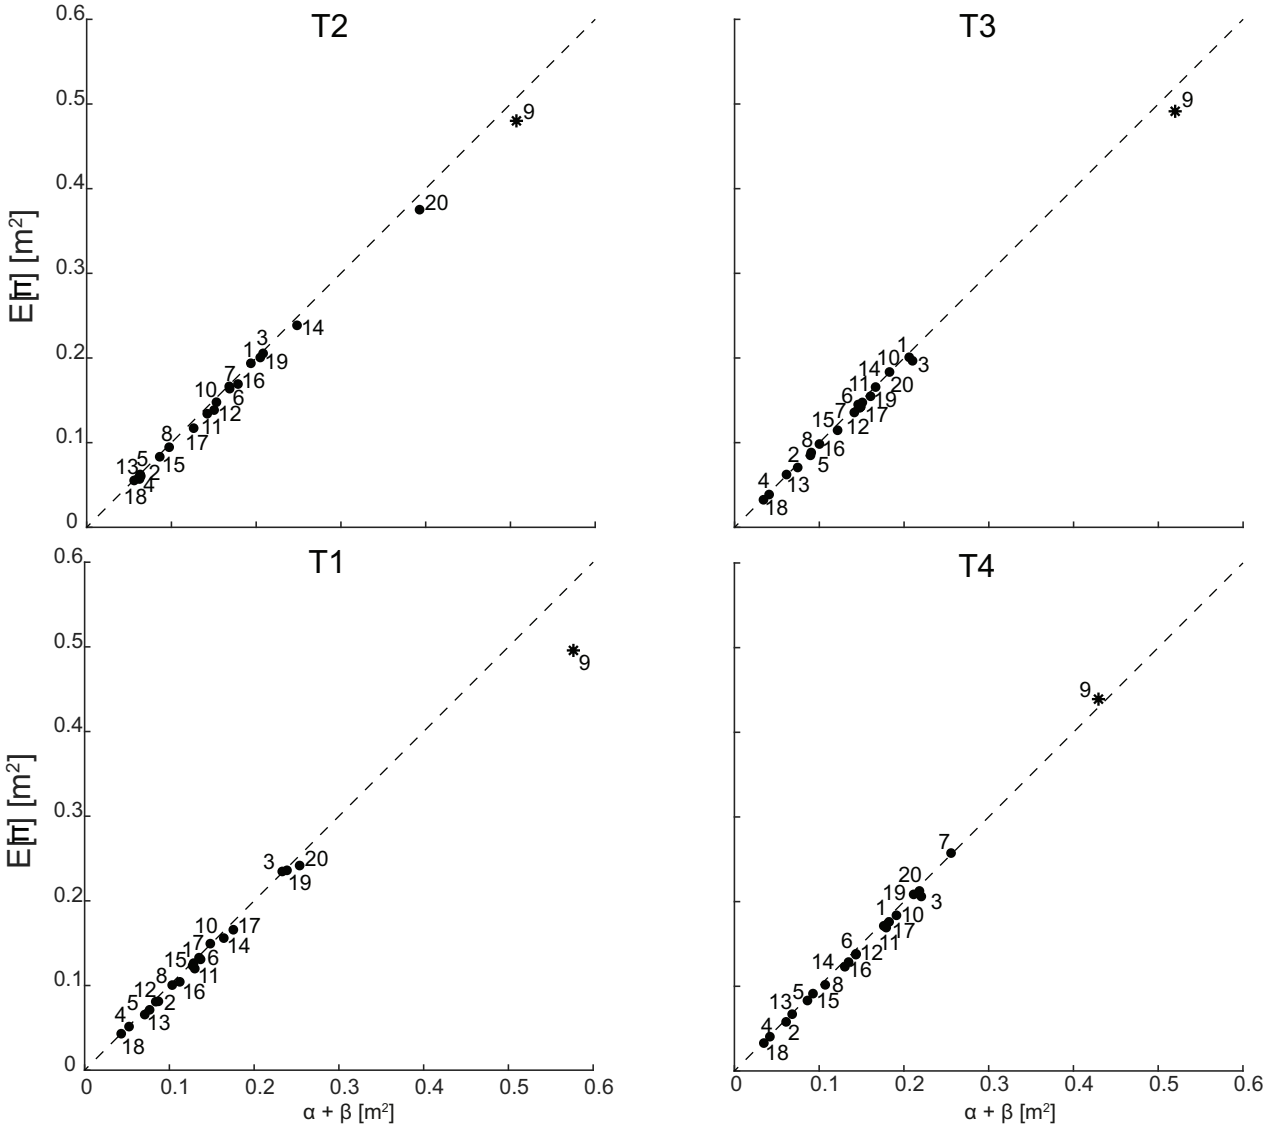

Figure 6. Validity of the quadratic approximation. (Sample) mean score vs local quadratic approximation (27) across targets (T1-T4, *different panels*) and participants ($P1$-$P20$) as in 1. Notice that $P9$ is the only participant is the only one deviating substantially from the identity line (*dashed line*).

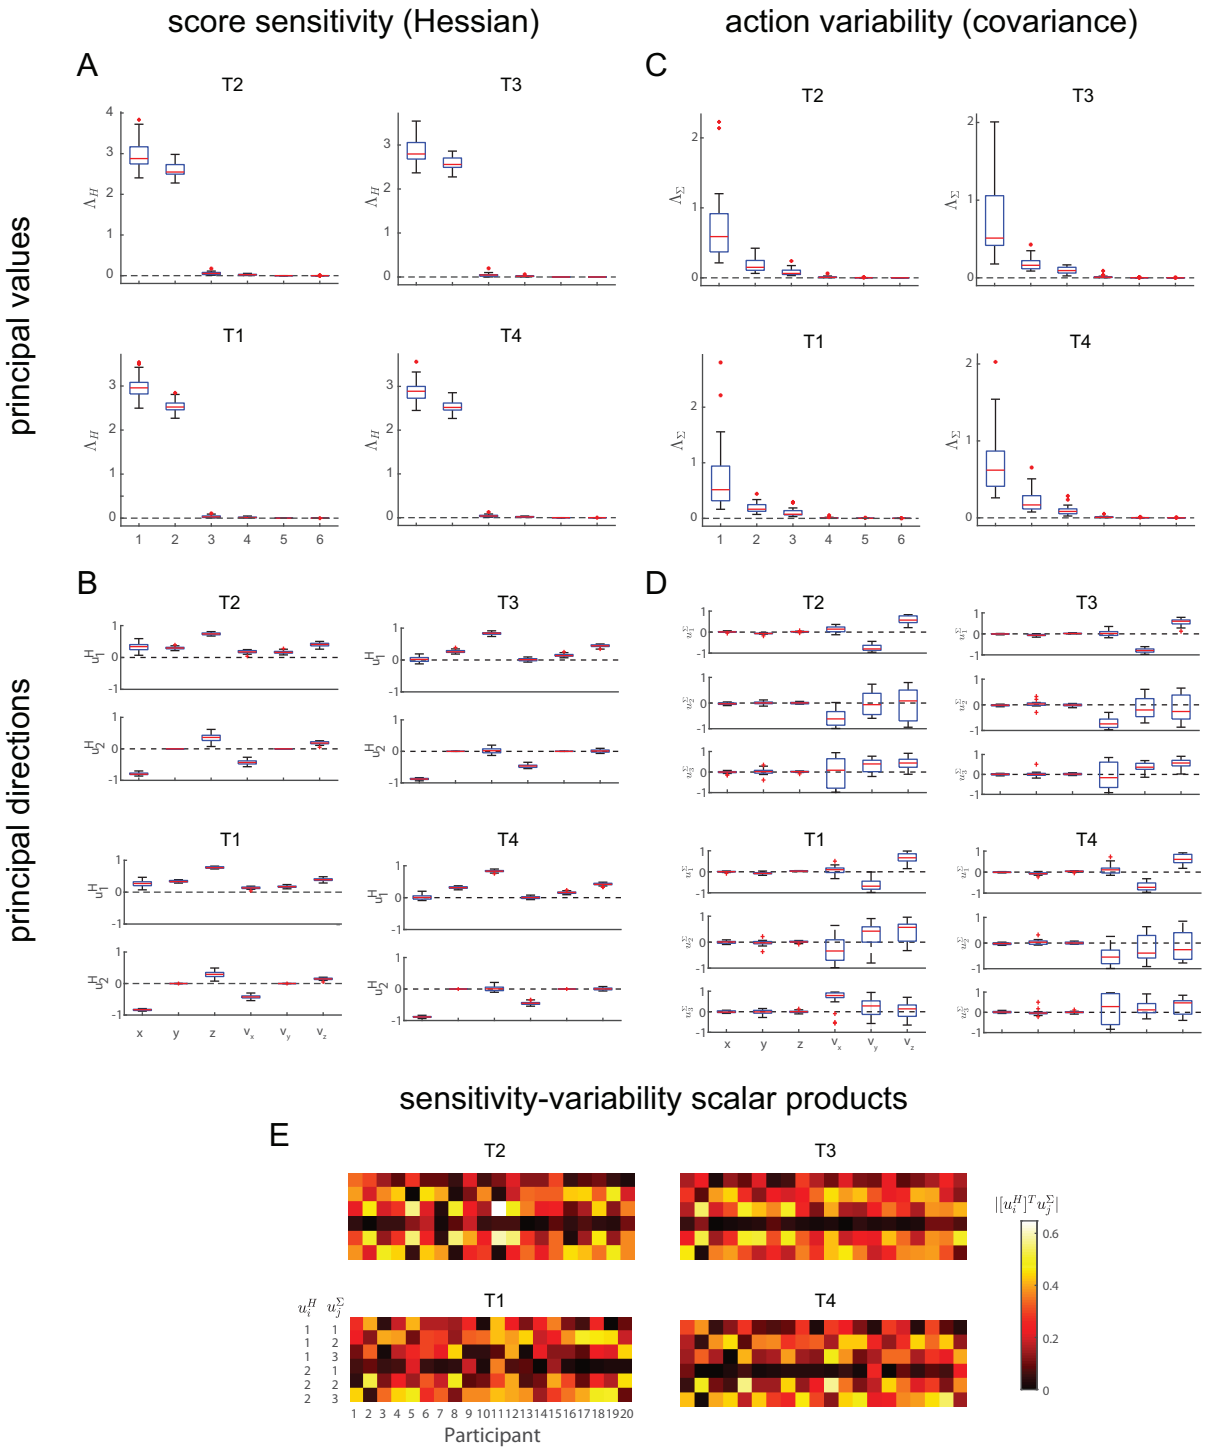
score sensitivity (Hessian)
action variability (covariance)
A
C
principal values
T2
T3
T1
T4
B
D
principal directions
x
y
z
v_x
v_y
v_z
sensitivity-variability scalar products
E
$|[u_i^H]^T u_j^\Sigma|$
u_i^H u_j^Σ
Participant

Figure 7. Sensitivity of the action-to-score function (Hessian matrix), action variability (covariance matrix), and their relationship across participants. (A) Distributions of the Hessian eigenvalues across participants for each target. In our scenario, the outcome space is bi-dimensional or, the solution manifold is a four-dimensional manifold embedded in the six-dimensional action space (see Appendix S1 Appendix C). Hence, the local tolerance of each participant is dominated by the first two eigenvalues of the Hessian matrix. (B) Distributions of the first ($\boldsymbol{u}_1^H$) and second ($\boldsymbol{u}_2^H$) principal sensitivity directions across participants for each target. The first principal sensitivity direction is dominated by the vertical release position and velocities while, the lateral and longitudinal components contributes 'equally' for target T1 and target T2, while for target T3 and target T4, the longitudinal components were 'more score relevant' than the lateral ones.. The second principal sensitivity direction is instead dominated by the lateral release position and velocity. (C) Distributions of the eigenvalues of the covariance matrix across participants for each target. The first three principal components can explain 95% of the total variance. (D) Distributions of the first ($\boldsymbol{u}_1^H$), second ($\boldsymbol{u}_2^H$), and third ($\boldsymbol{u}_2^H$) principal variability directions across participants for each target. (E) Absolute values of the scalar products between all pairs (i, j) of the first two principal sensitivity directions ($\boldsymbol{u}_i^H$) and the first three principal variability directions ($\boldsymbol{u}_i^\Sigma$).

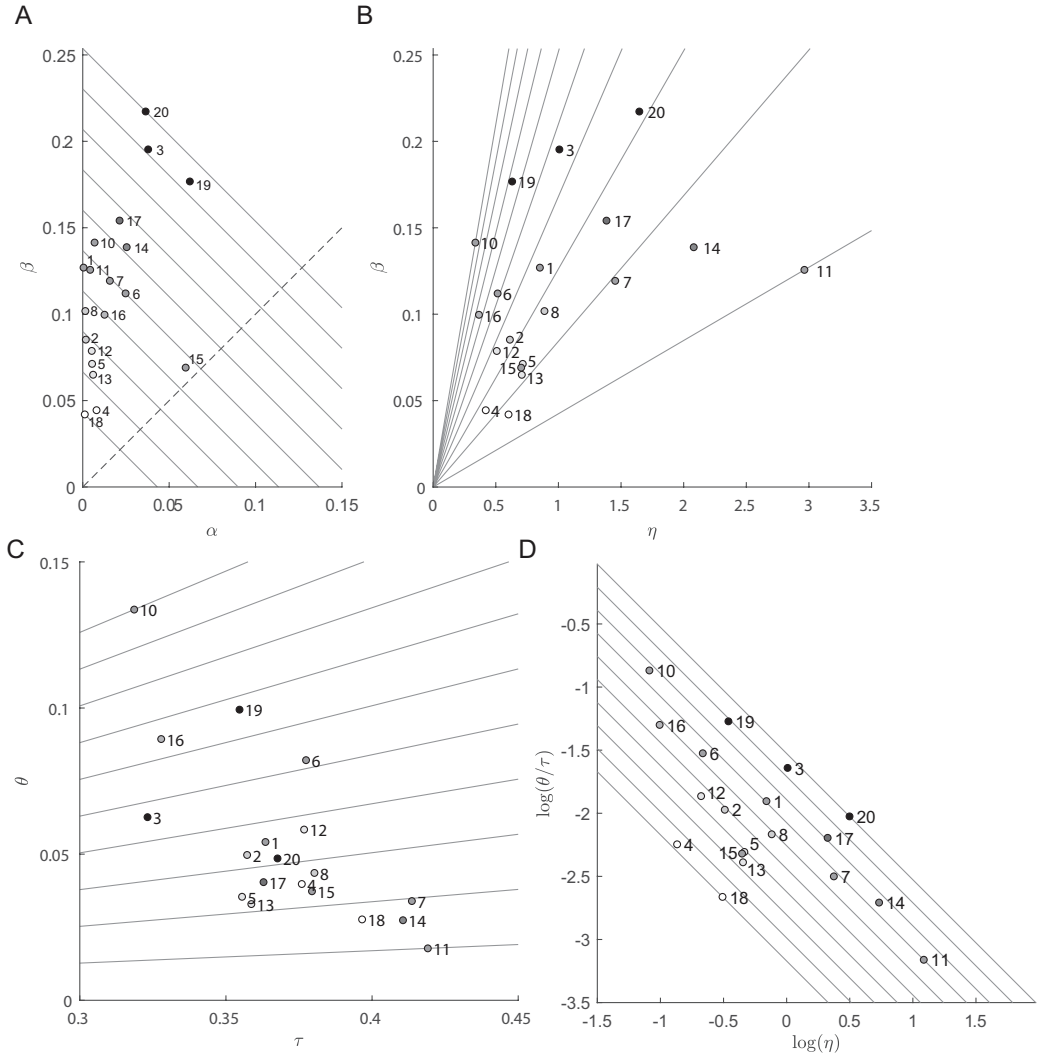

Figure 8. Decomposition of the mean score across participants for target T1. Markers are gray-shaded according to the score (lighter gray for lower mean score). To increase the readability of the figure, $P9$ is not shown. (A) The $\alpha - \beta$ plane and the iso-performance lines. The dashed line 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 uncorrelated noise η. (C) The $\theta - \tau$ (tolerance-alignment) plane. The lines have slopes corresponding to different values of $\frac{\beta}{\eta}$. 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.

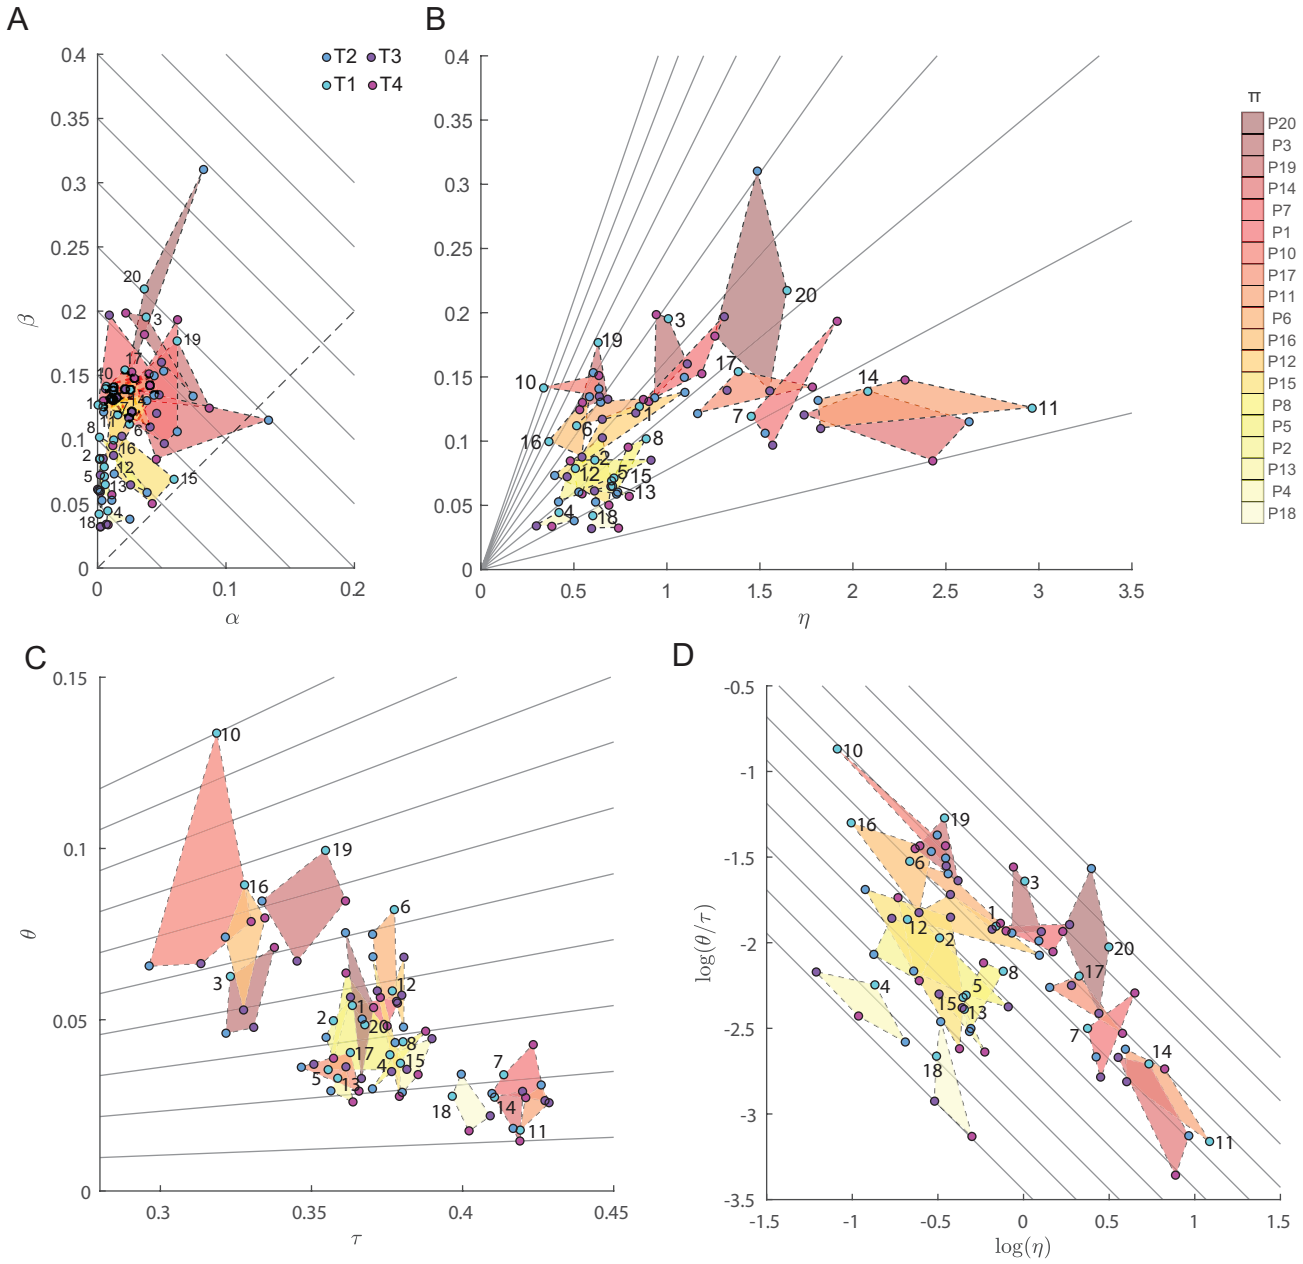

Figure 9. Decomposition of the mean score across participants for all targets. In each panel, the quadrilateral connects the points corresponding to the values of a pair of parameters for the four targets of each participant (excluding $P9$) and they are colored according to the mean score of that participant (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. (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.

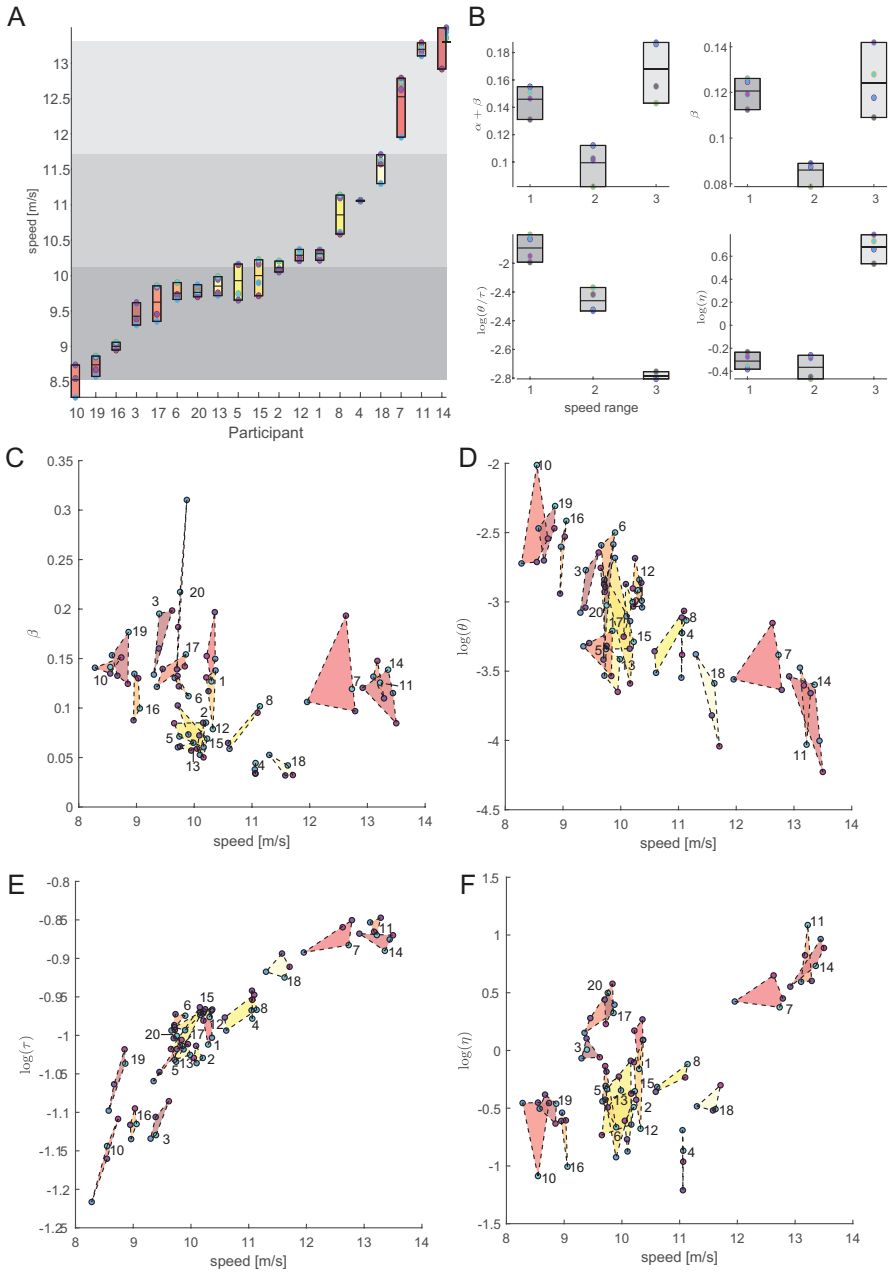

Figure 10. Dependence of decomposition parameters on throwing speed. (A) The distribution of the mean release speed of all participants (excluding $P9$), sorted according to mean speed, is shown for each target (*circular markers, colored as in Fig.9*) and as overall mean over targets (*horizontal line* within the *rectangle* representing the range of target values). (B) Mean value of $\alpha + \beta$ (*top left*), β (*top right*), $log(\theta)/log(\tau)$ (*bottom left*), and $log(\eta)$ (*bottom right*) across participants in three speed intervals. (C-F) Dependence of β (C), $log(\theta)$ (D), $log(\tau)$ (E), and $log(\eta)$ (F) on speed (same color coding of individual participants and targets as in Fig 9).

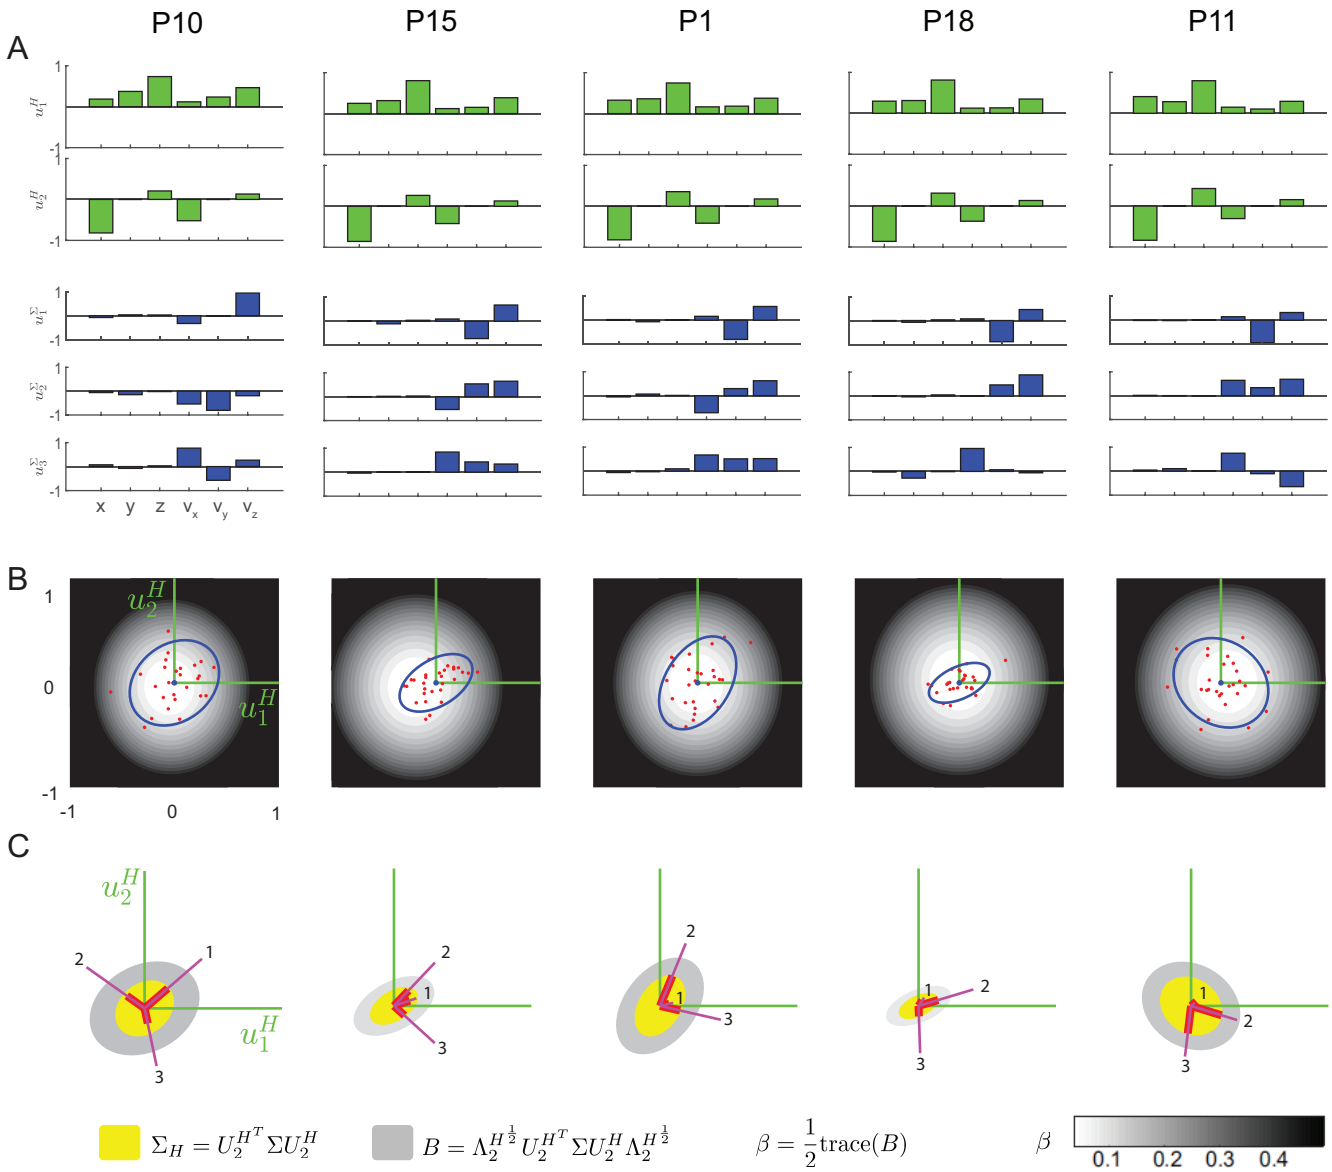

Figure 11. Low-dimensional representation of the β index. A: principal sensitivity directions (green bars) and principal variability directions (blue bars). B: Principal sensitivity plane: local action score (gray-shaded areas), first two principal sensitivity directions (green axes), mean action (large blue circles, at the origin of the axes), score-relevant perturbations (red small circles), and their covariance (1 standard deviation blue ellipses). C: Representation of the contributions of noise (η), alignment (θ), and tolerance (τ) to the tolerance-variability index (β). The unit length principal variability directions (thin magenta lines) and scaled by their principal values (thick red lines) are shown as projected on the principal sensitivity plane (green axes). The length of the three projected principal variability directions provides a visual representation of the alignment (θ) and indicates how much the total variation in the actions (η) contributes to the score-relevant variability, represented by one-standard deviation yellow covariance ellipses. The gray ellipses are instead a visual representation of the further contribution of score tolerance to the β index, as the local score sensitivity amplifies/attenuates the score-relevant variability (B matrix). Because β has the units of the score, the gray colormap gives an additional visual representation of the tolerance-variability index and facilitate the comparison across participants.

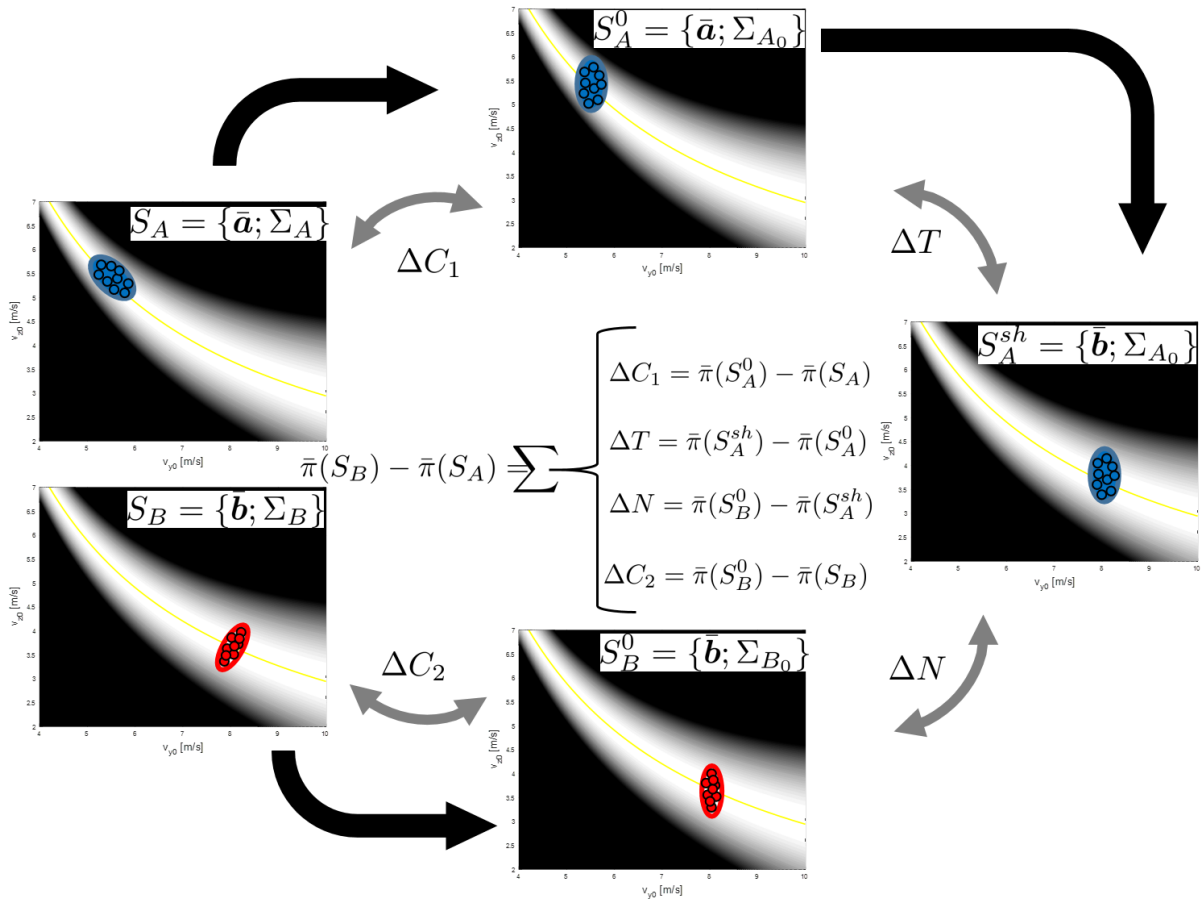

Figure 12. The TNC method proposed in [21]. Given two datasets, A and B, for instance the release strategies of two different partecipants, the method requires the generation of surrogate data-sets S^0 (covariation-free) and S^{sh} (covariation-free but shifted mean). The difference in mean score between experimental and surrogate datasets are then used to calculate the relative tolerance ΔT, noise ΔN and covariation ΔC_1, ΔC_2 between the two strategies.