# A Hessian-based decomposition to characterize 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
p.tommasino@hsantalucia.it (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 mapping and its geometric relationship with the action covariance determines the noise sensitivity of the action distribution. Thus, we approximate mean score as the sum of the score of the mean action and a tolerance-variability index which depends on both Hessian and covariance matrices. 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 our decomposition allows to compactly characterize inter-individual differences in throwing strategies with different but also with similar performance.

## 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 [14, 16, 11]. In this perspective, the goal of a thrower would be that of minimizing/maximizing the mean score over repeated trials [7, 19, 6]. Because of bias and noise in the sensorimotor transformations mapping goals into actions [10, 22], 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 [14, 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. 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 [14, 26].

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 [27, 29]. 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. First-order methods such as the Uncontrolled Manifold (UCM) [23, 15] and the Goal-Equivalent Manifold (GEM)[4, 5] 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 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 behavior 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, 26]. The Tolerance-Noise-Covariation (TNC) method 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 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 [26], 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 [18, 17], we found that throwing styles, i.e. the average trajectory of each limb segment, differed considerably across individuals. Similarly, we found important differences in individual performances, as shown in Fig 1E. 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. Fig 1C-E illustrates an example of two throwers, P1 and P11, achieving similar performances but characterized by different distributions of ball release position and velocity: the two throwers released the ball with different average position and velocity; furthermore, P1 had a larger variability in the ball release position with respect to P11, while the opposite held for the release velocity distributions.

Because of the need to describe unconstrained throwing action with, at least, six ball release parameters, we could not use the TNC method, since it requires a number of numerical operations that scale exponentially with the number of dimensions of the action space, to identify the different contributions of the action distribution to performance. We thus introduced a novel analytic method, based on second order derivatives of the action-to-score function, to identify how different characteristics of the action distribution contribute to performance. Our approach is based on the following assumptions: i) the action distribution is sufficienttly 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 map. Similarly to the TNC approach, our method allows to decompose the mean score of complex actions as the sum of the score of the mean action and a term that considers both action variability and action-to-score geometry. One advantage of our analytic approach with respect to TNC is that it can be applied to tasks involving high-dimensional actions. In addition, differently from TNC, the contributions to individual performances of motor variability, local geometry, and their alignment, can be estimated independently.

# 2 Methods

## 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 $\boldsymbol{a}_t \in R^n$ 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 an action score function $s^a : \boldsymbol{a} \in R^n \rightarrow R$. The score function can be seen as a cost, a reward, or a mixture of cost and reward functions, which punish and rewards 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 $\boldsymbol{x}_T$.

Since actions are stochastic, the action score $\pi$ will also be stochastic. Here we are interested in its *expected value* $E[\pi]$ as a measure of *individual (throwing) performance* and how it is affected by the skill-level and hence by the action distribution of an individual strategy.

More specifically, with reference to Fig 1A, we can say that an *individual release strategy is optimal* or *close to optimal* when *its expected action score* is zero, or close to zero ($E[\pi] \approx 0$), as for thrower $P18$ and $P4$ in Fig 1E. To understand why the individual release strategies of these two participants gets on average, less penalties, we have to evaluate the following integral:

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

which is the *expected action score* of an individual action distribution or strategy $p_{\boldsymbol{A}}$. Evaluating the above integral however, is not a simple task, especially when actions are high-dimensional ($n >> 1$) and the score function $s^a$ is non-linear. However, in what follows, we show that (1) can be approximated analytically for smooth score functions and knowing only the first and second moment of the action distribution, i.e. the mean action and the action covariance.

### The structure of an individual action strategy: principal components and noise $\eta$

Since in high dimensions it may be difficult and computationally expensive to infer the exact p.d.f. $p_{\boldsymbol{A}}$ of an individual action strategy, in the following we will 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 deifinite). 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 $\bar{\boldsymbol{a}}$. In summary, the mean action represents an individual preference in choosing, *on average*, a given action: for example releasing a ball, on average, from a certain position and with a given speed. Conversely, the covariance matrix $\Sigma^a = E[\boldsymbol{\delta_a}^T \boldsymbol{\delta_a}]$ represents action covariation/correlation (*action variability*) across multiple trials.

The action covariance matrix $\Sigma^a$, symmetric and positive-definite, can be decomposed into singular values and singular vectors:

$$\Sigma^a = U^{\Sigma} \Lambda^{\Sigma} U^{\Sigma^T} = \left[ \ \boldsymbol{u}_1^{\Sigma}, \boldsymbol{u}_2^{\Sigma}, ..., \boldsymbol{u}_n^{\Sigma} \ \right] diag(\lambda_1^{\Sigma}, \lambda_2^{\Sigma}, \ldots, \lambda_n^{\Sigma}) \left[ \ \boldsymbol{u}_1^{\Sigma}, \boldsymbol{u}_2^{\Sigma}, ..., \boldsymbol{u}_n^{\Sigma} \ \right]^T. \tag{3}$$

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

In the presence of motor redundancy, motor actions can be highly correlated across multiple trials. In such cases, the covariance matrix becomes almost singular and only $k < n$ singular values are significantly different from zero. Then, the principal variability directions $\boldsymbol{u}_i^{\Sigma}, i =$

$1 \dots k$ span a $k$-dimensional subspace of the action space and the covariance can be approximated with an $n \times n$ matrix of rank $k$:

$$\Sigma_k^a = U_k^\Sigma \Lambda_k^\Sigma U_k^{\Sigma^T} \tag{4}$$

where $U_k^\Sigma$ is an $n \times k$ matrix and $\Lambda_k^\Sigma$ 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.

In addition to the the principal components, and the principal variability directions, we will characterize an individual action strategy in terms of its *uncorrelated noise* (or total variation), which is defined as:

$$\eta = trace(\Sigma^a) = \sum_{j=1}^{n} \sigma_{jj} = \sum_{l=1}^{n} \lambda_l^\Sigma \tag{5}$$

### Quadratic approximation of the expected value of non-linear action scores

After choosing a model for the individual action strategy, in order to solve (1), we must find a suitable model for the high-dimensional action score, i.e. a trade-off between the complexity of the model and the accuracy with which we want to approximate the expected action score (1).

Compared to previous approaches, such as GEM and UCM, which have commonly assumed a linear or 'locally linear' relationship between performance score and actions, and hence between score variability and stochastic action perturbations, in this work we 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:

$$s^a(\bar{\boldsymbol{a}} + \boldsymbol{\delta_a}) \approx s^a(\bar{\boldsymbol{a}}) + [\boldsymbol{\nabla_a} s^s(\bar{\boldsymbol{a}})]^T \boldsymbol{\delta_a} + \frac{1}{2} \boldsymbol{\delta_a}^T H(\bar{\boldsymbol{a}}) \boldsymbol{\delta_a} \tag{6}$$

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

$$H_{ij}(\bar{\boldsymbol{a}}) = \frac{\partial^2 s^a}{\partial a_i \partial a_j}(\bar{\boldsymbol{a}}). \tag{7}$$

is the $n \times n$ symmetric Hessian matrix of the action score evaluated at $\bar{\boldsymbol{a}}$. Locally, i.e. around the mean action, the gradient represents the direction in action space where the score increases the most. It is constant and independent from the action only when the score is a linear function of the perturbation. When the score is non-linear and $\boldsymbol{\nabla_a} s^a(\bar{\boldsymbol{a}}) = 0$, then the average action corresponds to local minima/maxima of the score function. The Hessian $H(\bar{\boldsymbol{a}})$ instead, measures the local curvature, hence deviations from linearity, of the score function in a neighborhood of $\bar{\boldsymbol{a}}$. Notice that the Hessian matrix penalizes stochastic perturbations quadratically, meaning that 'the larger' the Hessian, or the local curvature of the score, the more action variability $\boldsymbol{\delta_a}$ affects the score. In other words the Hessian defines the local sensitivity or *intolerance* of the score function to stochastic perturbations $\boldsymbol{\delta_a}$ of the average action $\bar{\boldsymbol{a}}$.

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

$$E[\pi] \approx E[s^a(\bar{\boldsymbol{a}})] + [\boldsymbol{\nabla_a} s^s(\bar{\boldsymbol{a}})]^T E[\boldsymbol{\delta_a}] + E[\frac{1}{2} \boldsymbol{\delta_a}^T H(\bar{\boldsymbol{a}}) \boldsymbol{\delta_a}] \tag{8}$$

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 evalutaed 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 $trace(\frac{1}{2}H(\bar{\boldsymbol{a}})\Sigma^a)$ [9]. 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 $\Sigma^0 = 0$, or when $H$ and $\Sigma^a$ are 'orthogonal'. In all other cases this term will influence the expected action score.

In summary, assuming a locally quadratic score function, the expected action score (1) can be (locally) approximated as:

$$E[\pi] \approx s^a(\bar{\boldsymbol{a}}) + trace(\frac{1}{2}H^a\Sigma^a) = \alpha(\bar{\boldsymbol{a}}) + \beta(\bar{\boldsymbol{a}}, \Sigma^a) \tag{9}$$

where $\alpha(\bar{\boldsymbol{a}})$, is the score of the average action, and $\beta(\bar{\boldsymbol{a}}, \Sigma^a)$, the *tolerance-variability index*, is an index which captures how local non-linearities of the action score ($H(\bar{\boldsymbol{a}})$) and variability in the action strategy ($\Sigma^a$) affect (increase/decrease) performance, i.e. the average score $E[\pi]$.

## 2.2 The $\beta$ index and the interplay between score tolerance and action variability

Rewriting (9) as $\beta(\bar{\boldsymbol{a}}, \Sigma^a) \approx E[\pi] - \alpha(\bar{\boldsymbol{a}})$ shows that $\beta$ represents (approximately) the error which one would commit in estimating the expected action score $E[\pi]$ with a locally-linear assumption about the score function for which $E[\pi] \approx \alpha(\bar{\boldsymbol{a}})$, i.e. the error in assuming that expected action score is simply the score of the mean action and hence that non-linearities and action variability play no role on the mean action score, hence on performance.

This section will focus on the $\beta$ index and in particular on the interplay between score sensitivity/tolerance to stochastic perturbations and action variability. To ease the notation, we will write $H^a$ simply as $H$, however it must be kept in mind that the Hessian matrix will in general be a function of the mean action, and hence it may change across participants that have different action strategies. Similarly, individual action variability will be indicated simply as $\Sigma$.

### The local geometry of a score function and the tolerance $\tau$

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:

$$H = U^H \Lambda^H U^{H^T} = \left[ \boldsymbol{u}_1^h, \boldsymbol{u}_2^h, ..., \boldsymbol{u}_n^h \right] diag(\lambda_1^H, \lambda_2^H, \ldots, \lambda_n^H) \left[ \boldsymbol{u}_1^h, \boldsymbol{u}_2^h, ..., \boldsymbol{u}_n^h \right]^T. \tag{10}$$

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 $\left[ \boldsymbol{u}_1^h, \boldsymbol{u}_2^h, ..., \boldsymbol{u}_n^h \right]$. Singular values and singular vectors tell us many aspects about the local geometry/local curvature of the score function. For instance, when the Hessian matrix is positive-definite, i.e the singular values $\lambda_i^H$ are positive, than the mean action is in, or 'close to', a (local/global) minimum of the score function (the score is locally convex).

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. 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 $\lambda_i^H$ express the *sensitivity* of the score to stochastic perturbations, the larger $\lambda_i^H$, the more sensitive the score function is to perturbations $\boldsymbol{\delta_a}$ which are directed along the $i$-th singular 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 $trace(H) = \sum_{i=1}^{n} h_{ii}$, i.e. as the sum of the diagonal elements of the Hessian matrix. The *score tolerance* then, is defined as the inverse of the score sensitivity:

$$\tau = \frac{1}{trace(H)}. \tag{11}$$

**Score-relevant subspace and principal curvatures**

In many motor tasks, as in our throwing scenario, actions are redundant, i.e. there are many different actions that result in the same action outcome and hence in the same score.

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 sec. 4.3), Therefore, the $m$ *principal curvature 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 $m$:

$$H_m = U_m^H \Lambda_m^H U_m^{H^T} \quad \text{with } 1 \leq m < n \tag{12}$$

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

## 2.3 Decomposition of the tolerance-variability index $\beta$ and the alignment $\theta$

With the above definitions of tolerance and noise, by normalizing both the Hessian and the covariance matrix by their respective traces, we can rewrite $\beta = \text{trace}(\frac{1}{2}H\Sigma)$, the tolerance-variability index, 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* $\theta$:

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

is a scalar that measures the relative orientation between (normalized) principal curvatures and (normalized) principal components. 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 $\beta$ and hence on the expected action score.

### $\beta$ in the presence of redundant actions and the score-relevant variability

For symmetric and positive-definite Hessian matrices, $\beta$ can also be written as:

$$\beta = \text{trace}\left(\frac{1}{2}H^{\frac{1}{2}}\Sigma H^{\frac{1}{2}}\right) \tag{15}$$

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, than $\beta$ can also be approximated as:

$$\beta \approx \text{trace}\left(\frac{1}{2}\left(U_m^H \Lambda_m^{H\,\frac{1}{2}} U_m^{H^T}\right)\Sigma\left(U_m^H \Lambda_m^{H\,\frac{1}{2}} U_m^{H^T}\right)\right) \tag{16}$$

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

$$\Sigma_H = U_m^{H^T}\Sigma U_m^H = E\left[\left(U_m^H \boldsymbol{\delta_a}\right)^T\left(U_m^H \boldsymbol{\delta_a}\right)\right] \tag{17}$$

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

$$\beta \approx \text{trace}\left(\frac{1}{2}U_m^H \underbrace{\Lambda_m^{H\,\frac{1}{2}}\Sigma_H\Lambda_m^{H\,\frac{1}{2}}}_{B} U_m^{H^T}\right) = \frac{1}{2}\text{trace}(B) \tag{18}$$

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

## 2.4 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 [18] 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 [14, 24], 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 $\boldsymbol{x}_T$ [24]:

$$\pi = s^x(\boldsymbol{x}, \boldsymbol{x}_T) = (\boldsymbol{x} - \boldsymbol{x}_T)^T(\boldsymbol{x} - \boldsymbol{x}_T) \equiv e^2. \tag{19}$$

Written in this form, the accuracy score, represents a *task score* $s^x : R^2 \to R^+$ which penalizes the bi-dimensional action outcomes $\boldsymbol{x}$ with a scalar score $\pi$. To find the relationship between release actions and quadratic (task) error, i.e. the *action score function* $s^a : \boldsymbol{a} \in R^6 \to R^+$, we need to express the task score (19) 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:

$$\boldsymbol{f}(\boldsymbol{p}_0, \boldsymbol{v}_0, t) = \begin{cases} x(t) = x_0 + v_{x0}t \\ y(t) = y_0 + v_{y0}t \\ z(t) = z_0 + v_{z0}t - \frac{1}{2}gt^2 \end{cases} \tag{20}$$

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

For a target board oriented as in Fig 1A, i.e. with the normal pointing in the longitudinal direction $y$, the time of impact $T_i$ 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 (20)). Hence, at the time of impact, the projectile will hit the board at:

$$\boldsymbol{x}_i = \boldsymbol{f}(\boldsymbol{a}, T_i) = \begin{cases} x_0 + v_{x0}\frac{y_b - y_0}{v_{y0}} \\ y_b \\ z_0 + v_{z0}\frac{y_b - y_0}{v_{y0}} - \frac{1}{2}g\left(\frac{y_b - y_0}{v_{y0}}\right)^2 \end{cases} \tag{21}$$

Substituting the above system of equations into (19), let us writing the *action score* as:

$$e^2 = s^x(\boldsymbol{x}_i, \boldsymbol{x}_T) = s^x(\boldsymbol{f}(\boldsymbol{a}, T_i), \boldsymbol{x}_T) = s^a(\boldsymbol{a}, \boldsymbol{x}_T) \tag{22}$$

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

$$e^2 = s^a(\boldsymbol{a}, \boldsymbol{x}_T) = \left(x_0 + v_{x0}\frac{y_b - y_0}{v_{y0}} - x_T\right)^2 + \left(z_0 + v_{z0}\frac{y_b - y_0}{v_{y0}} - \frac{1}{2}g\left(\frac{y_b - y_0}{v_{y0}}\right)^2 - z_T\right)^2 \tag{23}$$

Given an individual release strategy with mean action $\bar{\boldsymbol{a}}_0$ and covariance $\Sigma^a$ and aiming at hitting a desired target $\boldsymbol{x}_T$, the mean squared error ($E[e^2] = E[\pi]$) can be approximated with (9):

$$E[\pi] = E[e^2] \approx \alpha(\bar{\boldsymbol{a}}) + \beta(\bar{\boldsymbol{a}}, \Sigma^a) = \alpha(\bar{\boldsymbol{a}}) + \frac{\theta(\bar{\boldsymbol{a}}, \Sigma^a)}{\tau(\bar{\boldsymbol{a}})}\eta(\Sigma^a) \tag{24}$$

where $\alpha(\bar{\boldsymbol{a}})$ is just (23) evaluated at $\bar{\boldsymbol{a}}$, i.e. the quadratic error of the outcome of the mean action, and $\beta$ can be decomposed into the three components $\eta$, $\tau$, and $\theta$ by using the covariance matrix of the action strategy $\Sigma^a$ and the $6 \times 6$ Hessian of (23) evaluated at $\bar{\boldsymbol{a}}$.

In this work, the $6 \times 6$ Hessian matrix of (23) is calculated with the MATLAB Symbolic Toolbox for each target condition and for each individual strategy.

### A simulated 2D example

Equation (24) 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* $\Sigma^a$ and the local curvature of the score $H(\bar{\boldsymbol{a}})$, which increase/decrease the $\beta$ 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 $\alpha$ and $\beta$ parameters close to zero. To give a visual representation of our approach Fig 2A shows a toy model of a dart throwing task. In this task, the action score (23) is assumed to be dependent only on the longitudinal and vertical release velocity, while (19) penalizes, quadratically, 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 the solution manifold, is represented with the red ellipses, whose major and minor axes are, the first and second principal curvatures 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 2B shows three simulated individual strategies, each drawing actions according to a bi-dimensional Gaussian distribution with mean action $\bar{\boldsymbol{a}}_i$ and covariance $\Sigma_i$, with $i = g, b, p$ ($g$ = green, $b$ = blue, $p$ = purple). Each colored circle represents an action drawn from each of the distribution and Fig 2B *center* and *right* show the ball trajectories of each action and the score distribution of each strategy, respectively. Fig 2C 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 also results into 'biased' throwing outcomes, whose average landing position does not coincide with the desired target position as shown by the green distribution in Fig 2B *center*. Notice, however, that while the $g$ strategy has the worse $\alpha$, it has the best (lowest) $\beta$ compared to $b$ and $m$. The $\beta$ score depends on the interplay between action variability and local curvature and hence on the three parameters $\eta, \tau$ and $\theta$ shown in Fig 2C. 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 $\beta$ are only due to the local tolerance $\tau$ and the alignment $\theta$ between principal curvatures and principal components. The blue and purple strategies, have a similar alignment $\theta$, given that, for both strategies, the principal component (direction of greatest variability) is almost parallel to the direction of maximal curvature. However, compared to the purple strategy, the blue strategy can afford a much lower $\beta$ because its mean action is located in a more tolerant region of the score, i.e. $\tau_b >> \tau_p$. While the 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 $\beta$, 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 curvature $\boldsymbol{u}_1^{H}$.

### 2.5 Experimental protocol and data analysis

Individual release strategies were obtained from the experimental dataset acquired in our previous study [18], where twenty right-handed participants (10 females, 10 males; age: $28.2 \pm 6.8$ 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 component 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 (21) and (19) 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 (6% of the total number of throws). The error distribution, across trials, participants and target conditions, between experimental and ideal performance (mean squared error) is shown in Fig 3 (mean $\pm$ SD: $0.0012 \pm 0.0912 m^2$). The dataset is available in S2 Dataset.

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 $\Sigma^a$, 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

As the Hessian-based performance analysis that we propose in this work 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 (23), we first assessed the validity of such assumption. In other words, we tested whether the score-relevant variability was not too large compared to the local tolerance and whether (24) can be considered as an acceptable model of the mean action score. Fig 4 shows the relationship (24) between the (sample) mean squared error and the sum $\alpha + \beta$ across participants and targets. We notice that the 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. It should be noted that this participant had the largest number of unsuccessful actions (balls not hitting the score board) that where removed from the estimation of its mean action and covariance. Hence the limited number of samples may have 'inflated' its covariance matrix and hence invalidating the assumption of *small* stochastic perturbations and non-linearities. Our analysis shows that $P9$ was in fact one of the most variable (large $\eta$) and least tolerant (small $\tau$) participants. To quantify the goodness of our second-order approximation of the sample mean score by the sum $\alpha + \beta$, we computed the fraction of variance accounted for (VAF) defined as:

$$\text{VAF} = \left[1 - \frac{\text{var}(E[\pi]_i - \alpha_i - \beta_i)}{\text{var}(E[\pi]_i)}\right] 100\% \tag{25}$$

which is 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. 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 $P9$.

## 3.1 Hessian-based decomposition of the mean score of individual strategies

Fig 5 shows the distributions of $\alpha$, $\beta$, $\eta$, $\tau$ and $\theta$ across participants and relative to target 1. To visualize the inter-individual differences in terms of these five parameters, we split the parameters into three subspaces: the $\alpha$-$\beta$ plane Fig 5A, the $\eta$-$\beta$ plane Fig 5B and the $\tau$-$\theta$ plane Fig 5C.

### The $\alpha-\beta$ plane and the trade-off between optimal mean action and optimal tolerance-variability

The $\alpha-\beta$ plane (Fig 5A) shows the performance of each individual thrower and its decomposition in terms of $\alpha$ and $\beta$. In this plane, (24) 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 $\beta$ and hence to the geometric relationship between action variability and score tolerance. In other words, compared to $P18$, $P1$ has a higher tolerance-variability index ($\beta$) and hence on average collects more penalties. Conversely, while $P5$ and $P15$ have the same $\beta$, hence the same contribution of tolerance-variability, $P15$ perform worse than $P5$ due to

a higher $\alpha$. 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 $\alpha$, 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 $\beta$ for higher $\alpha$ and vice versa. Fig 6 shows the action outcomes and their score distributions for four representative participants: $P1, P11, P15$ are iso-performing participants, $P1$ and $P11$ have $\alpha \approx 0$, their mean release action is optimal and does not influences the mean score. Conversely, $P15$, has a smaller $\beta$, hence the tolerance-variability score has less influence on the mean score, however this is compensated by a larger $\alpha$, hence a less optimal mean action and a large systematic error error between target position and average outcome. $P18$, who is the best thrower, with both small $\alpha$ and small $\beta$, hence outcomes which are unbiased with respect to the center of the target and precise, is shown for comparison.

### Noise, tolerance and alignment in non-expert throwers

Fig 5B shows the distribution across participants of the total variation, or uncorrelated noise $\eta$, and its relation with the tolerance-variability index $\beta$. According to equation (13), each participant lies on a straight line, passing through the origin ($\eta = 0, \beta = 0$) and with slope equal to the alignment-to-tolerance ratio $\frac{\theta}{\tau}$. Notice that there are many differences across participants. The best throwers such as $P8, P4, P13$ have relatively little action variability $\eta$, relatively little alignment-to-tolerance ratio and hence relatively little $\beta$ score. Conversely, less skilled throwers such as $P20, P3$, and $P19$ populate regions with an intermediate level of noise but relatively high $\beta$. 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 component and principal curvatures are aligned as shown in the $\tau - \theta$ plane Fig 5C. Intermediate and iso-performing throwers can be found all over the $\beta - \eta$ plane. Notice that despite $P11$ is about three times more variable than $P1$, they have similar $\beta$. 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 $\theta$). Lastly notice that participants that have the same or similar slopes, such as $P2, P12$ and $P1$ or $P18$ and $P14$, have the same alignment-to-tolerance ratio and therefore the differences in their $\beta$ score is only due to stochastic action variability $\eta$.

## 3.2 The local geometry of the action-score and the structure of individual action strategies

When actions are high-dimensional, as in our scenario, it is difficult to visualize both the score and the individual strategies in a single plot. Fig 7 shows the action score and individual release strategies just in terms of the release velocities and for five exemplary participants. Notice that participants have been sorted from top to bottom according to their local tolerance (see Fig 5C), hence $P10$ is the least tolerant and $P11$ the most tolerant. The gray-scale colors are the penalty associated to each release velocity according to (23). The domain of each 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 average action. The plots have 10 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 that has radius 0.2 $m$. Black-color actions instead, receive a score that is higher or equal than 1 and hence are actions which land at $1m$ from the center of the target. Notice that the score has different shapes (or geometry) across the three velocity planes and across participants. The wider the white areas around the mean release velocity (blue square), the more tolerant is the score to stochastic perturbations. For instance, in the $v_{x0} - v_{z0}$ plane, the score, locally looks like an ellipsoid, whose anisotropy suggests that the $v_{z0}$ direction is less tolerant/more sensitive to action variability. Also notice that the 'size' of the ellipsoid is subject-specific (as it depends on the individual mean release action) and hence, the large variability of $P11$, in the $v_{x0} - v_{z0}$ is partially compensated by a more tolerant (local) score.

Fig 7 also shows the individual release strategies in terms of mean release velocity (blue square) and velocity variability (two-standard deviation covariance ellipses). Notice that different participants have different release strategies, both in terms of the mean release velocity and in terms of the action variability. For instance, $P1$, on average, releases the ball with a faster vertical ($v_{z0}$) and slower longitudinal ($v_{y0}$) velocity, compared to other participants, such as $P11$, who instead, on average, throws with a fast longitudinal velocity and an approximately zero vertical velocity. Also notice different patterns of covariance/correlation across individual strategies, for instance, $P10$, who is the least variable participant (Fig 5B) shows no correlation between $v_{y0}$ and $v_{z0}$, while $P11$, who is the most variable participant (Fig 5B), shows a negative correlation: reducing the vertical release velocity proportionally to an increase in the longitudinal release velocity. This allows $P11$ to be less aligned with the principal curvature directions and hence to reduce the alignment-to-tolerance ratio. Notice that for $P15$, who has a non-zero $\alpha$, the mean release velocity is located closer to the edge of the white region and hence it is distant from the solution manifold.

### 3.2.1 Score-relevant dimensions and score-relevant variability affecting the $\beta$ index

The analysis of the structure of individual variability and its relation with the local geometry of the score by 'visual inspection' of subspaces of the action space is not feasible when the score is defined over a high-dimensional space of action variables. For high-dimensional and non-linear problems it may not be easy to understand what are the subspaces (of the action space) that are more relevant for the score and hence for performance. In our scenarios for instance, is the score more sensitive to variability in the release velocities or in the release positions, or in a combination of both?

In this section we show that, similar to the UCM and GEM method (that use Jacobian matrices to split motor variability along task-relevant and task-irrelevant directions), we can use the Hessian matrix to quantify *score-relevant* variability affecting the mean score.

Fig 8 shows the distributions, across participants and for each target, of the principal sensitivities (singular values) of the Hessian matrix. Because in our scenario the outcome space is bi-dimensional, the solution manifold is a four-dimensional surface embedded in the six-dimensional action space (see sec. 4.3). Hence, the Hessian matrix will only have two non-zero singular values, whose associated singular vectors defines a score-relevant plane. Away from the solution manifold however, the Hessian matrix is also influenced by the non-linearities introduced by the mapping $\boldsymbol{f}$ 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 8 shows, however, that the contribution of the third and fourth singular value is negligible compared to the first two.

Focusing on the first two singular values, we notice that the sensitivity is slightly anisotropic, with the first principal curvature, on average, about 10% higher than the second. Furthermore, the first sensitivity shows the largest variability across participants and target conditions.

The two principal curvature 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 9 shows the distributions of the principal curvature direction components across participants and for each target. We notice that, across targets, the first principal curvature direction $u_1^H$ is dominated by the vertical components (both position and velocity) of the release parameters, while, the lateral and longitudinal components contribute 'equally' for Target 1 and Target 2, while for Target 3 and Target 4, the longitudinal components are 'more score relevant' than the lateral ones. This asymmetry is due to the fact that Target 1 and Target 2 had a larger lateral displacement with respect to starting position of our throwers (see Experimental Protocol and Fig 1A), and hence, the lateral release velocity becomes *relevant* to successfully hit these two targets.

The second principal curvature direction is instead dominated by the lateral components ($x_0$, $v_{x0}$) of the release parameters. Taken together, these results highlight that the squared error of throwing outcomes is slightly more sensitive to throwing variability directed along the sagittal plane than the frontal plane. This is not surprising given that the vertical trajectory of the ball, is influenced, non-linearly, by the gravity field.

In terms of action variability, for all participants and for all target conditions three principal components were able to explain 95% of the total variance, as shown in Fig 10. Notice that, across participants and target conditions, the eigenvalue of the first principal component shows large variability across participants. It should be noted that the number of principal components is coordinate dependent and in our scenario the action vectors contains 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 [13].

The bar plots in Fig 11A shows the individual principal curvatures and principal variability directions for five representative participants. Notice that, as also shown in Fig 9, there is little difference in terms of principal curvature directions (blue 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 (black 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 $\beta$ index. However, because in our scenario the Hessian is locally dominated by two principal curvature 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 curvature plane in a neighborhood of the mean release action (blue square) of each participant and the score-relevant perturbations of the mean release action

(yellow squares). To represent the curvature plane and action score in Fig 11B, we plot the first and second principal curvature 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 'tangential' (to the curvature plane) perturbation of the average action $\bar{\boldsymbol{a}}$, and hence its score can be calculated with (23) using $\boldsymbol{a}_0 = \bar{\boldsymbol{a}_0} + \boldsymbol{p}$. In this way it is possible to visualize the best (local) planar representation of the score (gray-level). The greay-level code for the action score is the same as in previous figures, and therefore actions inside the white disk land inside the 20 *cm* target radius and hence the optimal action resulting in zero penalty scores lies at the center of the white disk. Notice that the 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 curvature frame (blue lines) for $P15$ has a large shift compared to the center of white, 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 white disk in this representation).

Fig 11C shows how the individual principal components of variability contributes to the score-relevant variability and how the local sensitivity amplifies its detrimental effect for performance. The black thin lines are the principal variability directions (eigenvectors of the action covariance matrix) projected on the principal curvature plane, i.e. $U_2^H U_3^\Sigma$ and therefore are a visual representation of the 'alignment' between principal curvature and principal variability direction. By rescaling each eigenvector by their respective eigenvalues, i.e. $U_2^H U_3^\Sigma \Lambda_3^{\Sigma^{\frac{1}{2}}}$, the black thick lines show the contribution of each principal component of variability to the score relevant variability. Notice that, on the principal curvature plane, $P10$, which is one the most aligned participant, has the largest projection of the principal variability directions, especially the first, when compared to least aligned participants such as $P18$ and $P11$.

The score-relevant variability (17) and the $B$ matrix in (18) are represented as a 2D ellipses with unitary standard deviation parameter [30] . The gray ellipses show the effect of the tolerance, and hence how the local sensitivity amplifies the score-relevant variability (yellow ellipse) of each participant. Notice that $P10$, $P1$ and $P11$ have a similar score-relevant variability (area of the yellow ellipses), however, $P10$ has a larger (and darker) ellipse due to a larger sensitivity of the score.

# 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 distance of the outcome of an action from the goal can be assessed by a score and that the score is a smooth function of the action, we derived an approximate but analytical relationship between the mean score and the first two moments of the actions distribution: the mean action $\bar{\boldsymbol{a}}$ and action covariance $\Sigma^a$ 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 $\alpha$ parameter, when different from zero, measures deviations of the mean action from the set of actions that accurately achieve the goal (solution manifold). The $\beta$ 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: the *total action variability* ($\eta$), computed as the sum of the variances of the individual components of the action vector (i.e. the trace of the action covariance matrix); the *tolerance* of the action-to-score function ($\tau$), 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); the *alignment* ($\theta$), 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 $\theta$ parameters (Fig 5). In line with previous works focusing on low-dimensional throwing tasks [21, 4], in our unconstrained high-dimensional throwing task we found that skilled participants, such as $P18$ and $P4$, have small $\alpha$ (accurate mean action) and small $\beta$ (tolerance-variability index). Still, it is possible to differentiate two different optimizing strategies, as the low $\beta$ in $P4$ is achieved by minimizing action variability, while in $P18$ by compensating the higher variability in action execution ($\eta$) with higher tolerance $\tau$ and smaller alignment (i.e., smaller $\theta$). Different combinations of bias, noise, tolerance, and alignment, resulting in similar performances, are even more prominent in less skilled participants. For example, participants $P1$, $P11$ and $P15$ have similar performance but different proportions of $\alpha$ and $\beta$. Participant $P15$ trades off a high $\alpha$ with a low $\beta$ while the opposite occurs for participants $P1$ and $P11$. Moreover, participants $P1$ and $P11$ have similar bias and similar tolerance-variability index values but different levels of noise, tolerance, and alignment (Fig 5 and 6). Participant $P11$ has three times more noise ($\eta$) than participant $P1$ but an alignment-tolerance ratio ($\theta/\tau$) three times smaller because of both larger tolerance and smaller alignment. Such interplay between variability 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 7) 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.

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 [23, 15, 21, 12], 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 [31], 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 Eq. (8) and Eq. (9), 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, 27, 26]. 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. Sternad and collaborators have shown that the changes in action variability across two practice sessions (each a series of trials) can be quantified by computing the difference in performance, defined as the mean of the minimum distance of the ball from the target, as the sum of three components [21]. The *covariation* component represents the difference in mean score due to the different amount of covariation between the actions of each series. It is computed, independently for each session, by comparing the actual mean score with the one corresponding to a surrogate data distribution generated by randomly permuting each components of the action vectors across different trials (thus having same mean but zero covariance). It is important to notice that this term depends not only by the action covariances, but also by the local sensitivity of the action-to-score map in correspondence of the mean actions. The *task tolerance* component corresponds to the difference in mean score associated with the zero covariance surrogate distributions, due to the different location of the mean action of the two series. Finally, the *noise* component measures the remaining difference in mean score due to the different variability (noise) of the individual components of the two series, once the differences in mean action and covariation have been accounted for.

The TNC methods takes 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 method concerning the use of numerical procedures for generating surrogate data, which limits its applicability to high-dimensional actions. Moreover, our method allows to

decompose the contribution to performance of individual action distributions rather than the differences between pairs of distributions, and to determine the contribution of the local geometry and of the action variability independently on each other. As detailed in the Appendix (see sec. 4.1), 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. Thus, the three TNC terms, not only refer to differences between action distributions rather than to specific features of each distribution as the terms of our decomposition, but they also do not identify and isolate clearly the different features of the action distribution and of the action-to-score mapping geometry that determine performance. Furthermore, as shown in Appendix (see sec. 4.1), the TNC terms are asymmetric with respect to the moments of the two action distributions, which reflect the fact that the result of the decomposition depends on the sequence of application of the three steps [21]. 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.

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 4, 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.

Our decomposition method relies on the computation of the action covariance $\Sigma^a$ 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 $\eta$ and the tolerance $\tau$, being defined as traces of the covariance and Hessian matrices, respectively, change under coordinate transformations (unless a metric is chosen [2]). However, the $\alpha$ 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 $\beta$ term is the trace of the product of the covariance and the Hessian matrices and it is invariant under affine coordinate transformations, given that $\Sigma^a$ and $H^a$ transform in opposite ways (see Appendix 4.2). Thus, re-scaling of positional and velocity coordinates due to different choices of measurement units do not affect the decomposition of mean score as a sum of $\alpha$ and $\beta$. However, $\beta$ 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 [25]. 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 [28]. In our decomposition, if the mean of the action distribution is on the solution manifold ($\alpha = 0$), $\beta$ 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. $\eta$ is small) the change in $\beta$ due to non-linear coordinate transformations may be negligible.

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 [24]; and a model-free system in which learning is driven by reinforcement and punishment of successful/erroneous actions [8, 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 $\alpha$ and $\beta$. Do inter-individual differences in terms of $\alpha$ and $\beta$ translate into individual differences in terms of performance improvement? In future works 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.

# Appendix

## 4.1 Relation to Muller & Sternad 2004

In this section we show that, when the score function is smooth, and actions are linearly correlated across muiltiple trials, it is possible to derive (Hessian-based) analytic expressions to isolate the three components of the TNC approach [21], shown in Fig12. Given two experimental strategies, such as $S_A$ and $S_B$ in the figure, the 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: $\Delta C_1$, 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 'delta' in performance $\Delta C_2$ 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. $\Delta N = \bar{\pi}(S_B^0) - \bar{\pi}(S_A^{sh})$.

When motor strategies are drawn from a symmetric distribution, and motor actions are linearly correlated, $S_A = \{\bar{\boldsymbol{a}}; \Sigma_A\}$ 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 $\Sigma_{A_0}$ and $\Sigma_{B_0}$ of the uncorrelated strategies $S_A^0$ and $S_B^0$, can be simply computed as $\mathrm{diag}(\Sigma_A)$ and $\mathrm{diag}(\Sigma_B)$, i.e. as the matrix of the diagonal elements (variances) of $\Sigma_A$ and $\Sigma_B$, 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 $\Delta C_1$,$\Delta T$, $\Delta N$ and $\Delta C_2$, simply as:

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

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

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

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

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

$$\begin{aligned} \Delta\bar{\pi} = \bar{\pi}(S_B) - \bar{\pi}(S_A) = \Delta C_1 + \Delta T + \Delta N + \Delta C_2 \approx \\ \bar{\pi}(\bar{\boldsymbol{b}}) - \bar{\pi}(\bar{\boldsymbol{a}}) + trace(H_{\bar{\boldsymbol{b}}}\Sigma_B) - trace(H_{\bar{\boldsymbol{a}}}\Sigma_A) = \Delta\alpha + \Delta\beta \end{aligned} \tag{26}$$

## 4.2 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" [25]. 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:

$$\boldsymbol{b} = \boldsymbol{g}(\boldsymbol{a}) \tag{27}$$

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

$$\boldsymbol{\delta}^b = J(\bar{\boldsymbol{a}})(\boldsymbol{\delta}^a) \tag{28}$$

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

The score $\pi$ 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:

$$\pi = s^b(\boldsymbol{b}) = s^b(\boldsymbol{g}(\boldsymbol{a})) = s^a(\boldsymbol{a}) \tag{29}$$

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:

$$\frac{\partial s^a}{\partial \boldsymbol{a}} = J^T \frac{\partial s^b}{\partial \boldsymbol{b}} \tag{30}$$

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

$$\frac{\partial^2 s^a}{\partial \boldsymbol{a}^2} = J^T \frac{\partial^2 s^b}{\partial \boldsymbol{b}^2} J + \left(\frac{\partial s^b}{\partial \boldsymbol{b}}\right)^T \frac{\partial J}{\partial \boldsymbol{a}} \tag{31}$$

Hence:

$$H^a = J^T H^b J + \frac{\partial J}{\partial \boldsymbol{a}}^T \boldsymbol{\nabla_b} s^b \tag{32}$$

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 $\beta$ does not depend on the co-ordinates used to parameterise the action space. In fact, let $\Sigma^b$, 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:

$$\Sigma^a = J^{-1} \Sigma^b J^{-1^T} \tag{33}$$

and hence:

$$\beta = trace(\frac{1}{2} H^a \Sigma^a) = trace(\frac{1}{2} J^T H^b J J^{-1^T} \Sigma^b J) = trace(\frac{1}{2} H^b \Sigma^b) \tag{34}$$

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 (32) may not be negligible. In such case, the Hessian looses its tensorial property, and $\beta$ becomes a co-ordinate dependent measure:

$$\beta = trace(\frac{1}{2} H^a \Sigma^a) = trace(\frac{1}{2} H^b \Sigma^b) + trace\left(\frac{1}{2}\left(\boldsymbol{\nabla_b} s^b\right)^T \frac{\partial J}{\partial \boldsymbol{a}} \Sigma^b\right) \tag{35}$$

### 4.3 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 (32) becomes:

$$H^a = J^T H^x J + \frac{\partial J}{\partial \boldsymbol{a}}^T \boldsymbol{\nabla_x} s^x \tag{36}$$

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

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

6. Gregory Dam, Konrad Kording, and Kunlin Wei. Credit Assignment during Movement Reinforcement Learning. *PLoS ONE*, 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. Adrian M. Haith and John W. Krakauer. Model-based and model-free mechanisms of human motor learning. *Advances in Experimental Medicine and Biology*, 2013.

9. 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.

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

11. 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.

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

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

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

15. 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.

16. Harry Manley, Peter Dayan, and Jörn Diedrichsen. When money is not enough: Awareness, success, and variability in motor learning. *PLoS ONE*, 2014.

17. 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.

18. 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.

19. David Marc Anton Mehler, Alexandra Reichenbach, Julius Klein, and Jörn Diedrichsen. Minimizing endpoint variability through reinforcement learning during reaching movements involving shoulder, elbow and wrist. *PLoS ONE*, 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. Leslie C. Osborne, Stephen G. Lisberger, and William Bialek. A sensory source for motor variation. *Nature*, 2005.

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

24. 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.

25. 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.

26. 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.

27. 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.

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

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

30. G. Vallabha. *Plot Gaussian Ellipsoids, MATLAB Central File Exchange*, 2007.

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

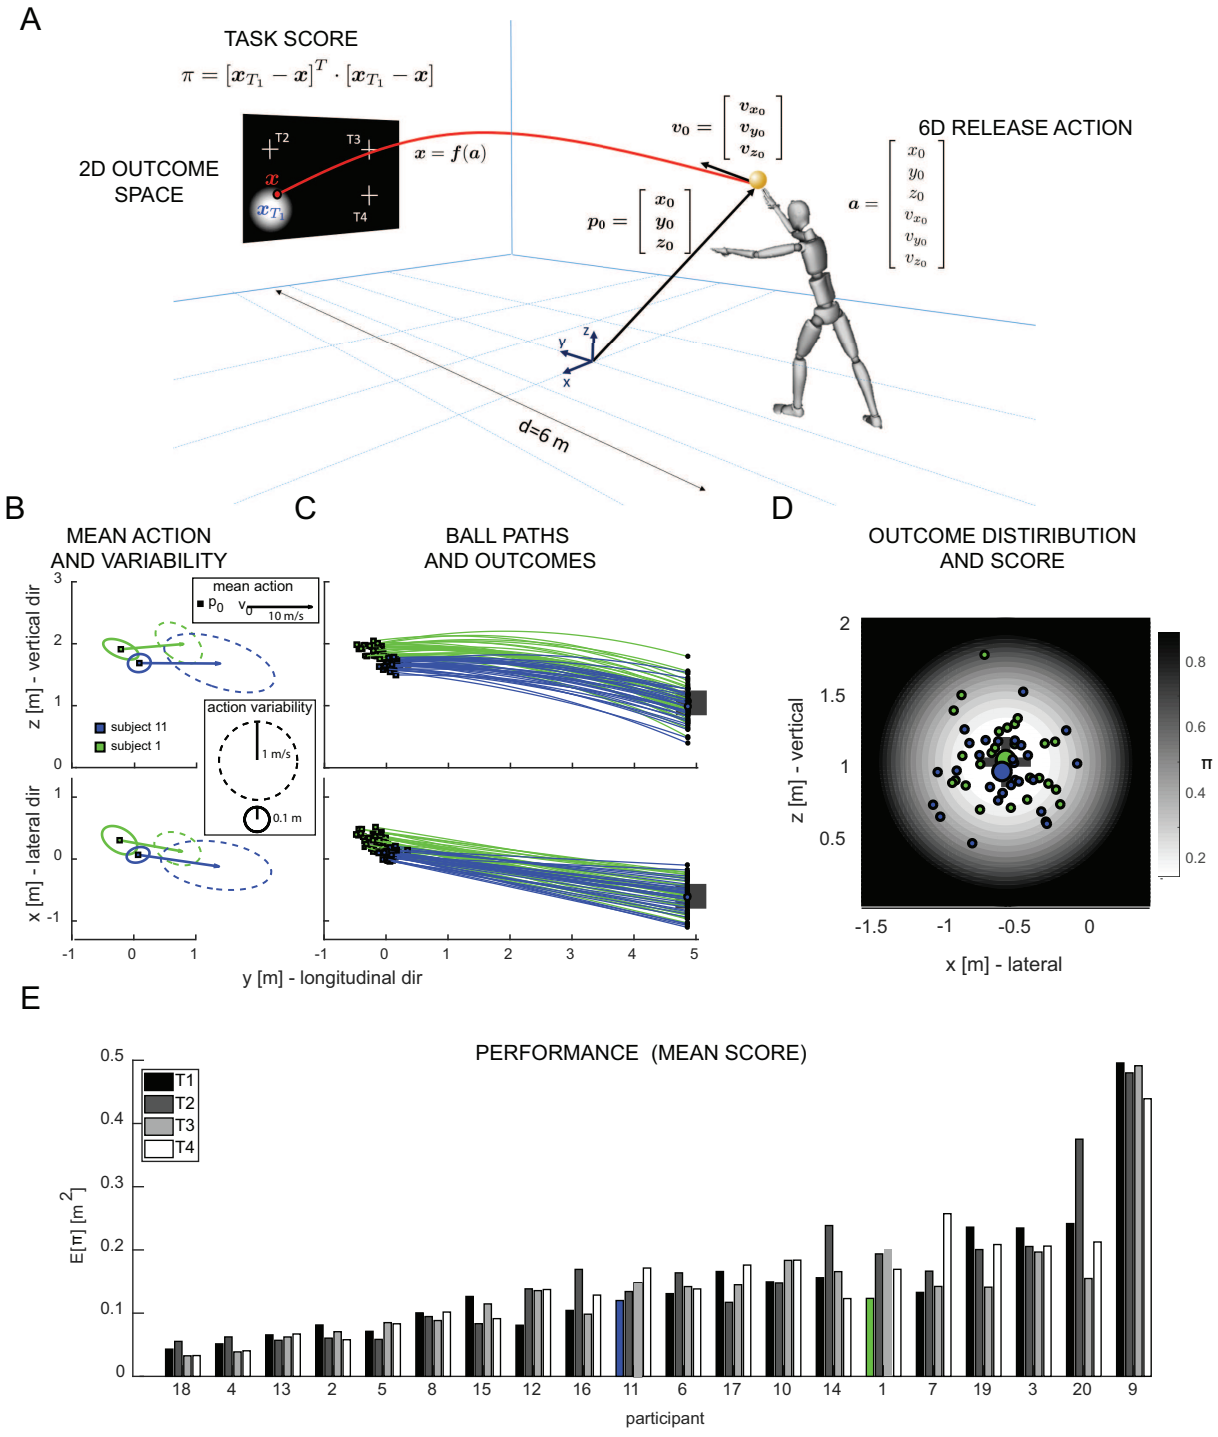


**Figure 1.** Throwing task and performance. (A) Schematic representation of our unconstrained overarm throwing scenario. (B) Example of two individual release strategies across multiple trials. Markers and arrows represent the mean release position and velocity respectively. While ellipses are a schematic representation of variability, across multiple trials, in the release position and velocity. (C) Ball paths and action outcomes of each individual strategy. (D) Action outcomes and $\pi$ score (squared error from the target center). (E) Individual performance or mean squared error. Notice the blue and green bar: despite the two participants have different release strategies, they have a very similar performance.

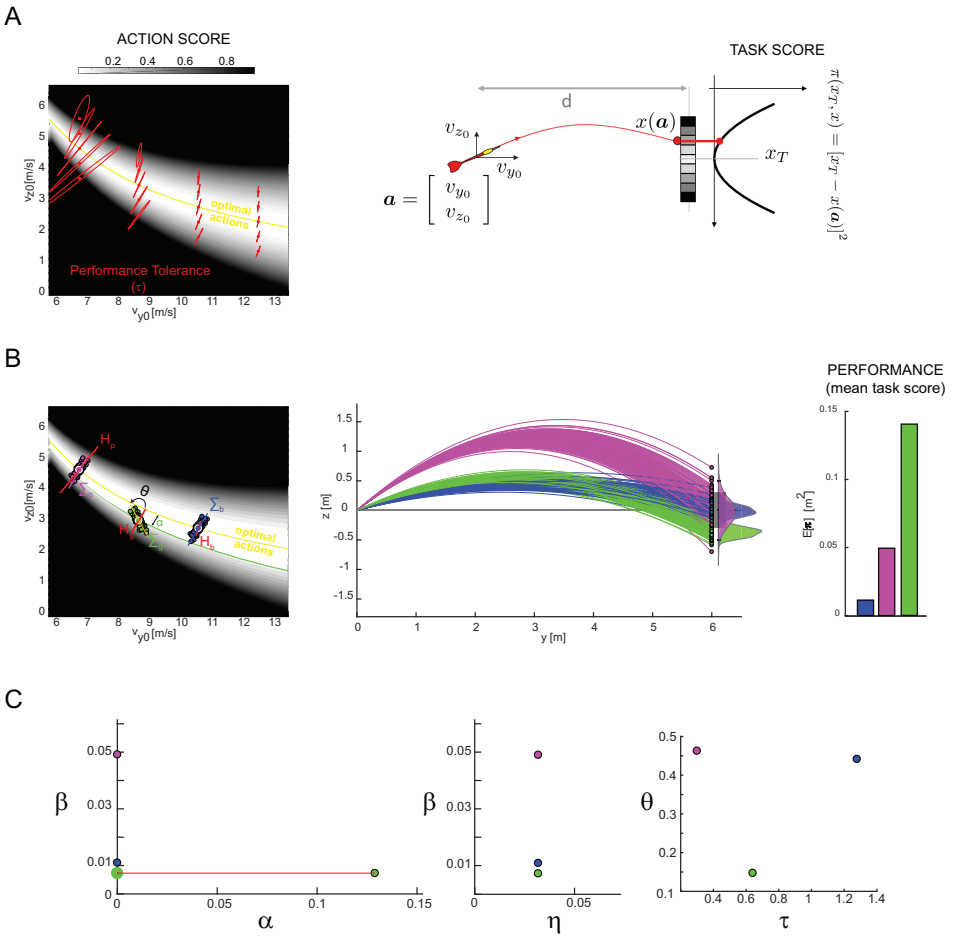


**Figure 2.** Simulated example of three different throwing strategies in a 2D throwing task. (A) Toy model of a throwing task involving a quadratic task-score. The action score and the solution manifold (yellow line) are shown in the left panel. The score sensitivity (the Hessian) is represented with 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 (24). The blue and purple strategies have zero $\alpha$ (score of mean action) as the mean of the distribution of release parameters is on the solution manifold but they have different $\beta$; the green strategy has a non-zero $\alpha$ corresponding to a non-optimal mean action. Notice that the release strategies have been chosen to have the same level of noise $\eta$ so as to highlight the effect of $\alpha$, $\tau$ (tolerance) and $\theta$ (alignment) on the individual expected score. See the text for more details.

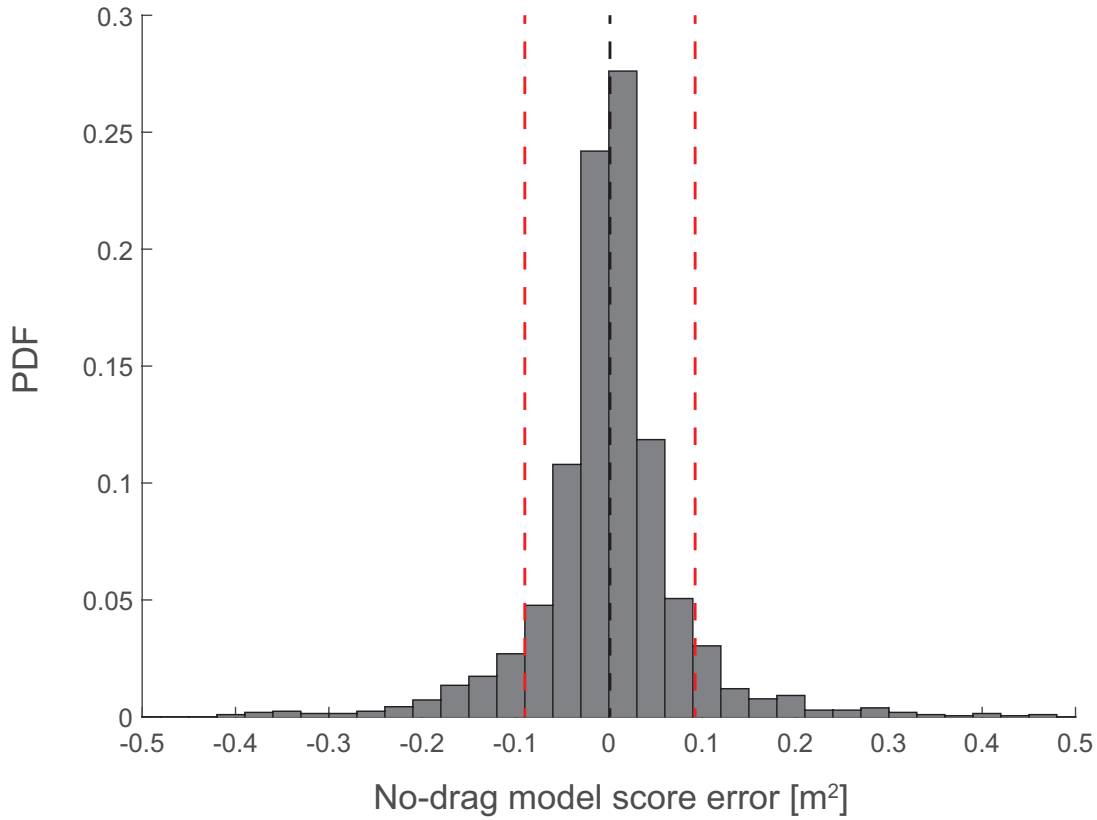


**Figure 3.** Error distribution between experimental performance and those predicted with the no-drag model in (23). Black and red lines indicate mean and $\pm 1SD$ of the distribution, respectively.

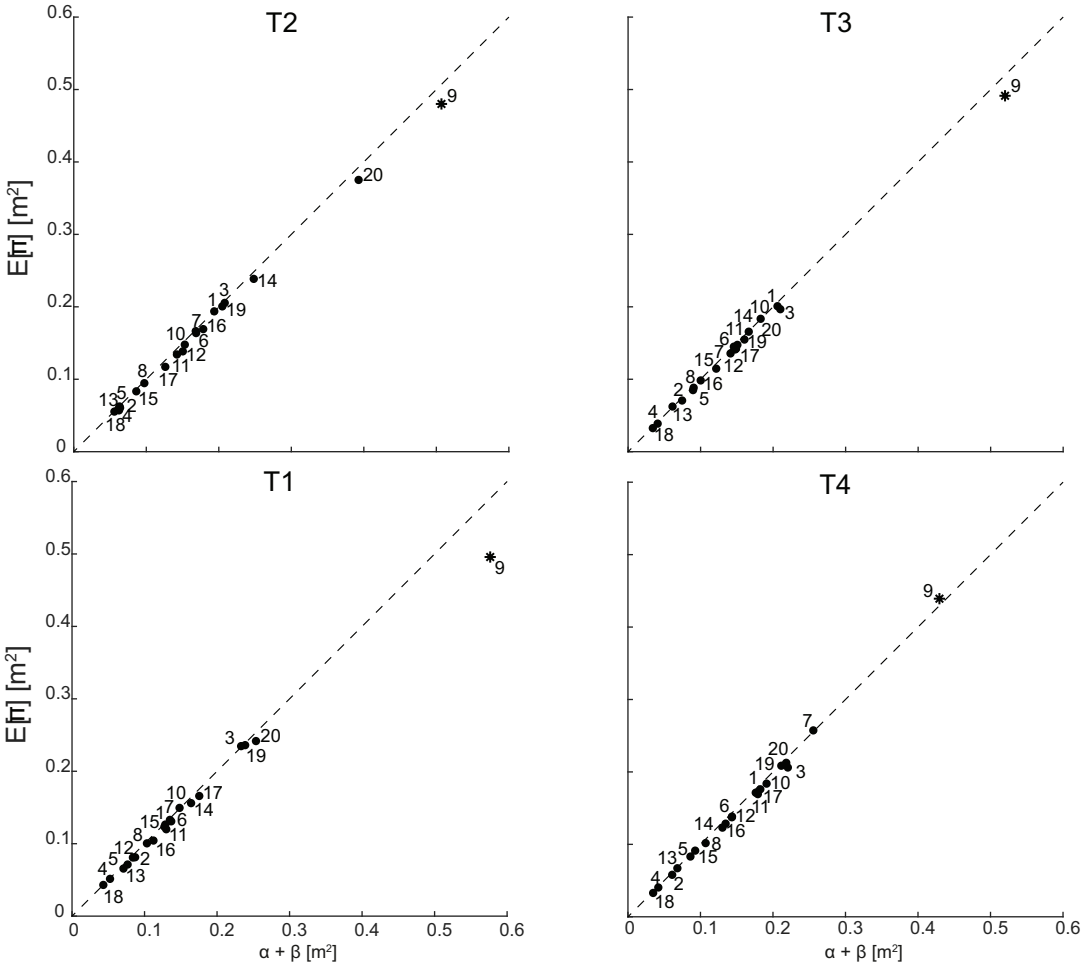


**Figure 4.** Validity of the quadratic approximation. (Sample) mean score vs local quadratic approximation (24) across targets (T1-T4, *different panles*) and participants (1-20) as in 1. Notice that the outlier $P9$ was not included in the linear regression (*dashed line*).

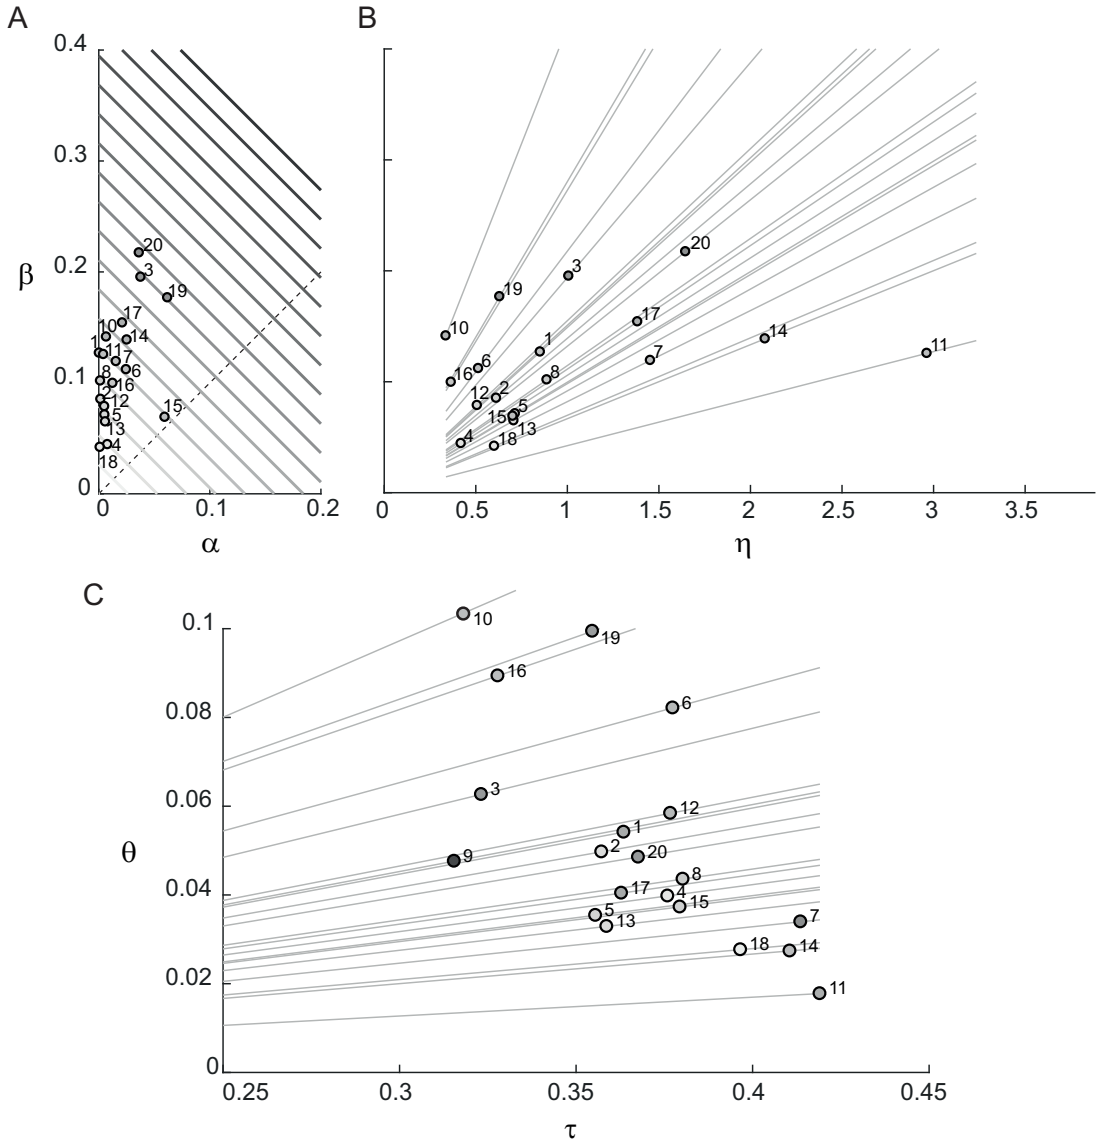


**Figure 5.** Decomposition of the mean score across participants for target 1. (A) The $\alpha - \beta$ plane and the iso-performance lines. Notice that to increase the readability of the figure, $P9$ is not shown as it has the largest $\alpha$ and $\beta$. (B) The $\beta - \eta$ plane. Markers are gray-shaded according to the expected score (lighter gray for lower mean score). The lines are drawn for each participant and have slope $\frac{\sigma_H}{\tau}$. In other words, participants such as $P1$, $P2$, and $P12$ have the similar alignment-to-tolerance ratio and the difference in their $\beta$ are only due to the uncorrelated noise $\eta$. (C) The $\theta - \tau$ (tolerance-alignment) plane.

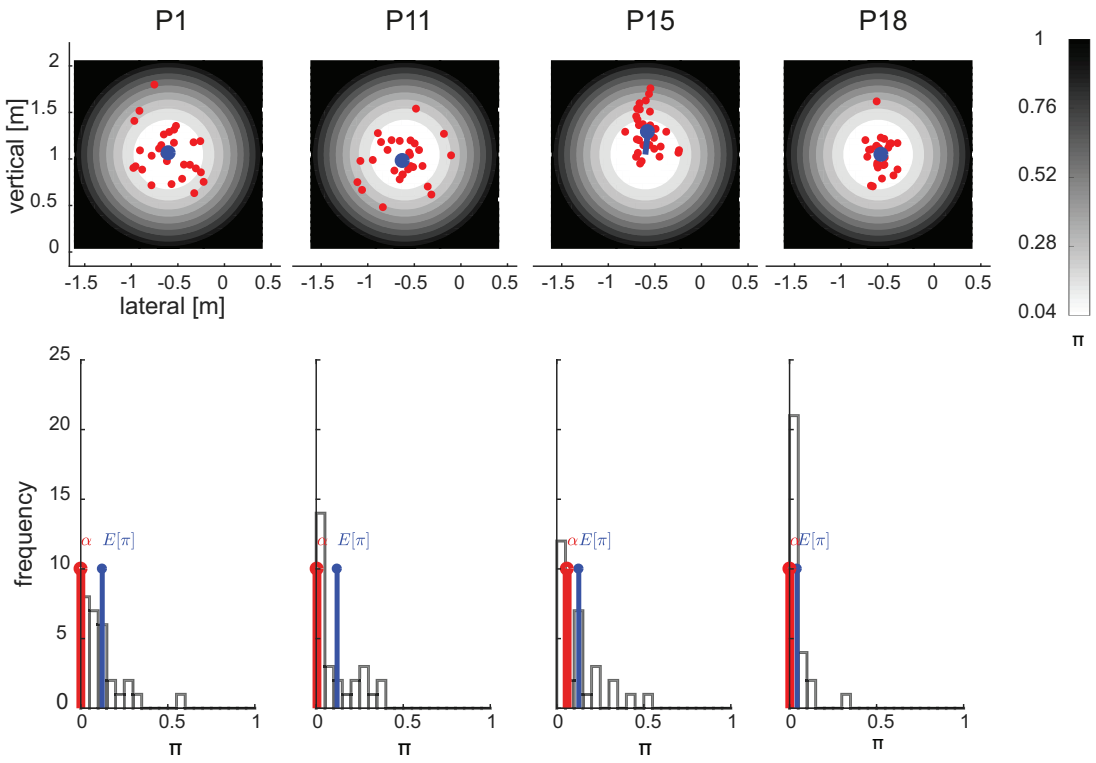


**Figure 6.** Examples of distributions of throwing outcomes and scores. *Top panel*: throwing outcomes (red circles) relative to Target 1 for different representative participants. The blue circle marks the mean outcome and hence the systematic error of each throwing strategy. The score ($\pi$) is represented with a gray-scale map and with 10 different gray-scale levels representative of (19). Balls landing within the white disk are balls which hit the target, hence that receive a score that is always smaller than 0.04 (the radius squared). *Bottom panel*: score (squared-error) distribution, its expected value $E[\pi]$ and the $\alpha$ parameter for each participant. The gap $E[\pi] - \alpha$ represents the $\beta$ parameter. Notice that to larger $\alpha$ there corresponds a larger systematic error ($P15$) and that to smaller $\beta$ corresponds more precise task outcomes ($P4$ and $P15$).

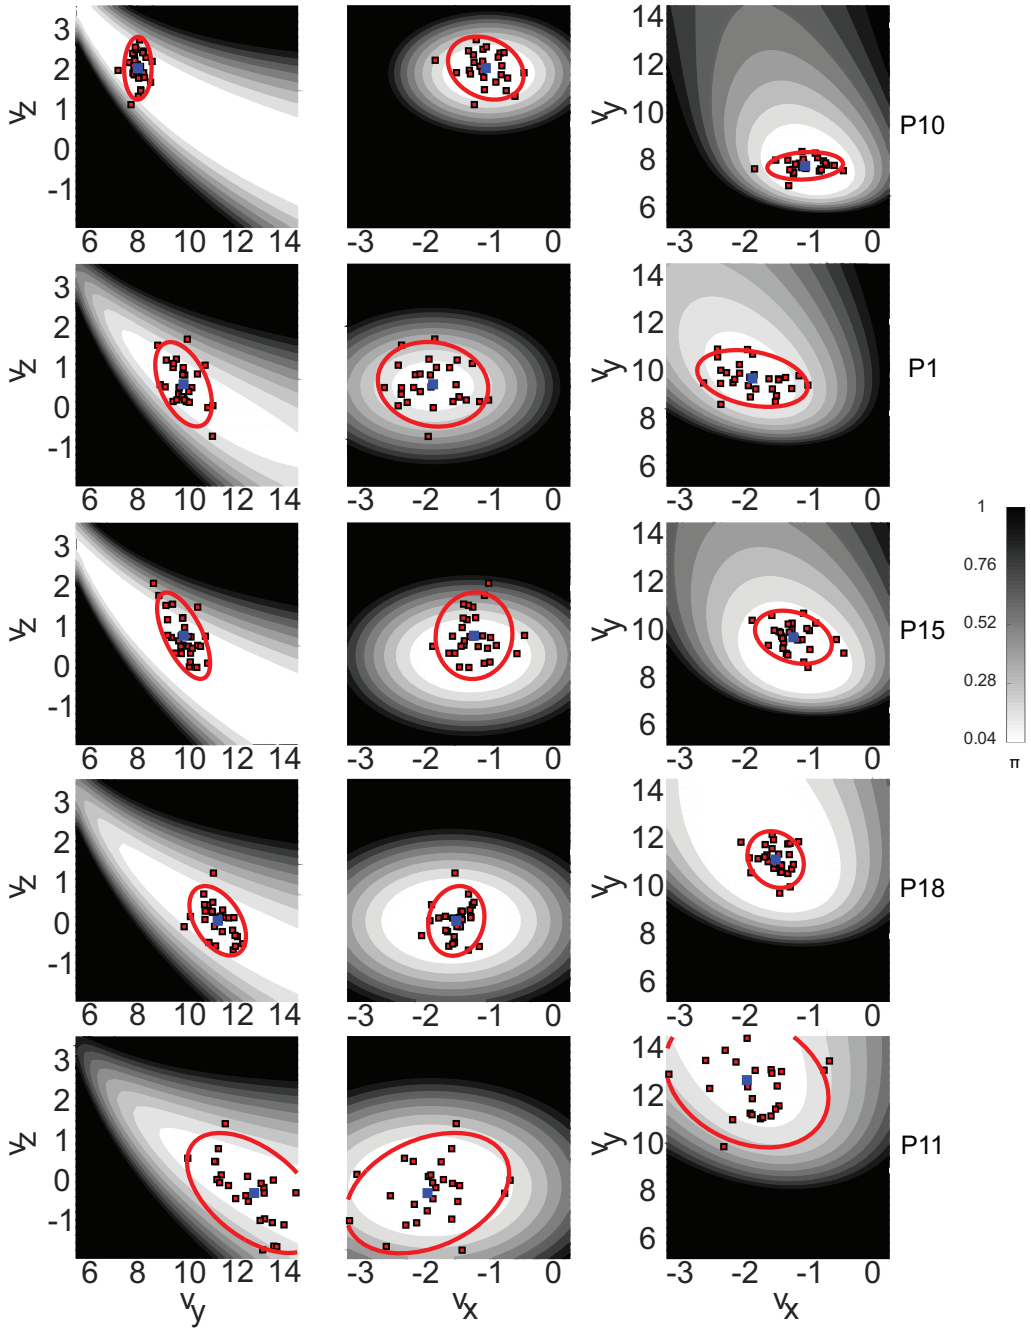


**Figure 7.** Examples of distribution of release velocities for five representative participants (target 1). Red squares represent release velocities of different throws. Blue squares and red ellipses represent means and covariance ellipses (two-standard deviations) of each strategy. The gray level map shows the local score, as a function of the release velocity, underlying each strategy. The wider the white area around the mean action, the more tolerant the score is to stochastic perturbations. Participants have been sorted from top to bottom according to their local tolerance (see Fig5C)

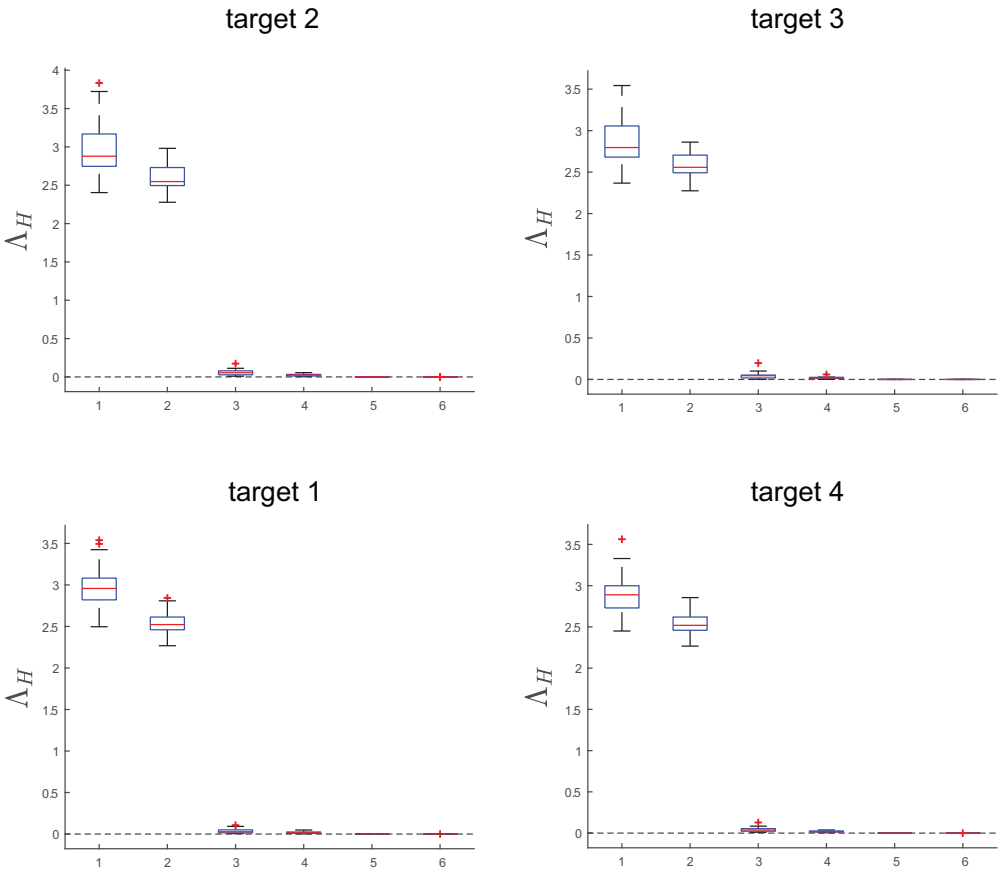


**Figure 8.** Distributions of the Hessian eigenvalues across participants for each target.In our scenario, the outcome space is bi-dimensional or, the solution manifold is a four-dimensional surface embedded in the six-dimensional action space (see sec. 4.3). Hence, the local tolerance of each participant is dominated by the first two eigenvalues of the Hessian matrix.

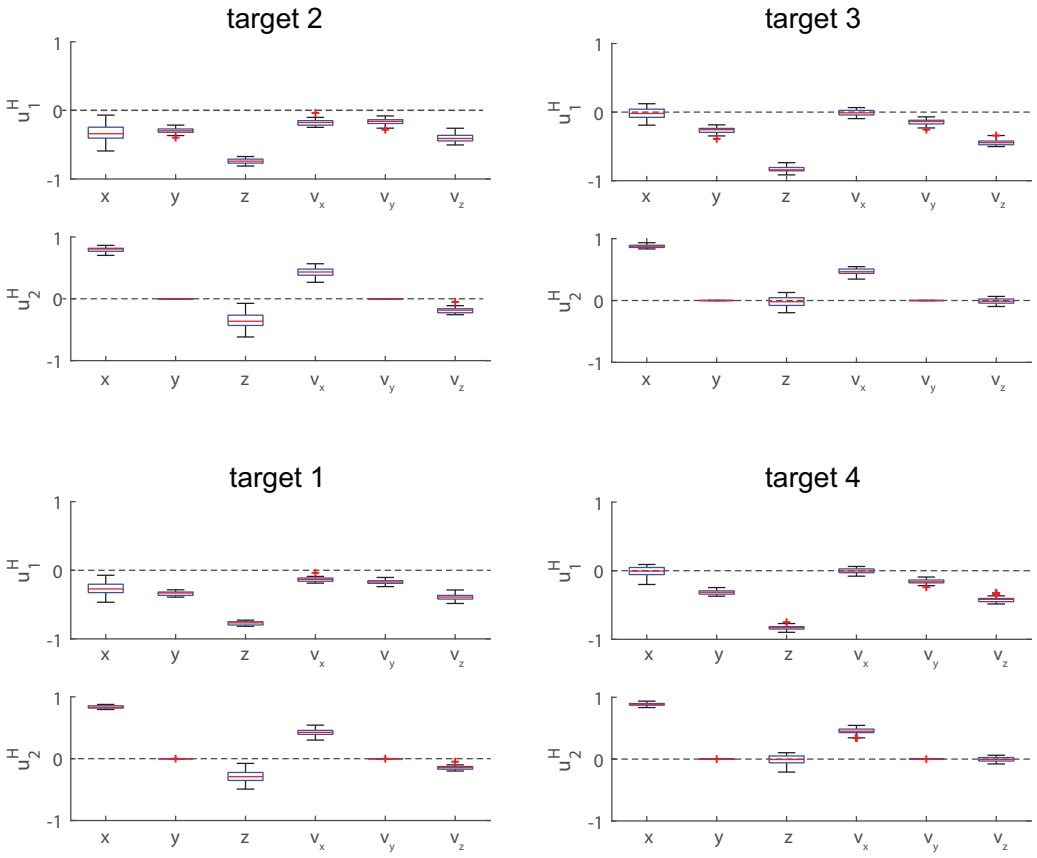


**Figure 9.** Distributions of the first ($u_1^H$) and second ($u_2^H$) principal curvature directions across participants for each target. The first principal curvature direction is dominated by the vertical release position and velocities while, the lateral and longitudinal components contributes 'equally' for Target 1 and Target 2, while for Target 3 and Target 4, the longitudinal components were 'more score relevant' than the lateral ones.. The second principal curvature direction is instead dominated by the lateral release position and velocity. Taken together, these results highlights that the squared error of throwing outcomes is slightly more sensitive to throwing variability directed along the sagittal plane than the frontal plane.

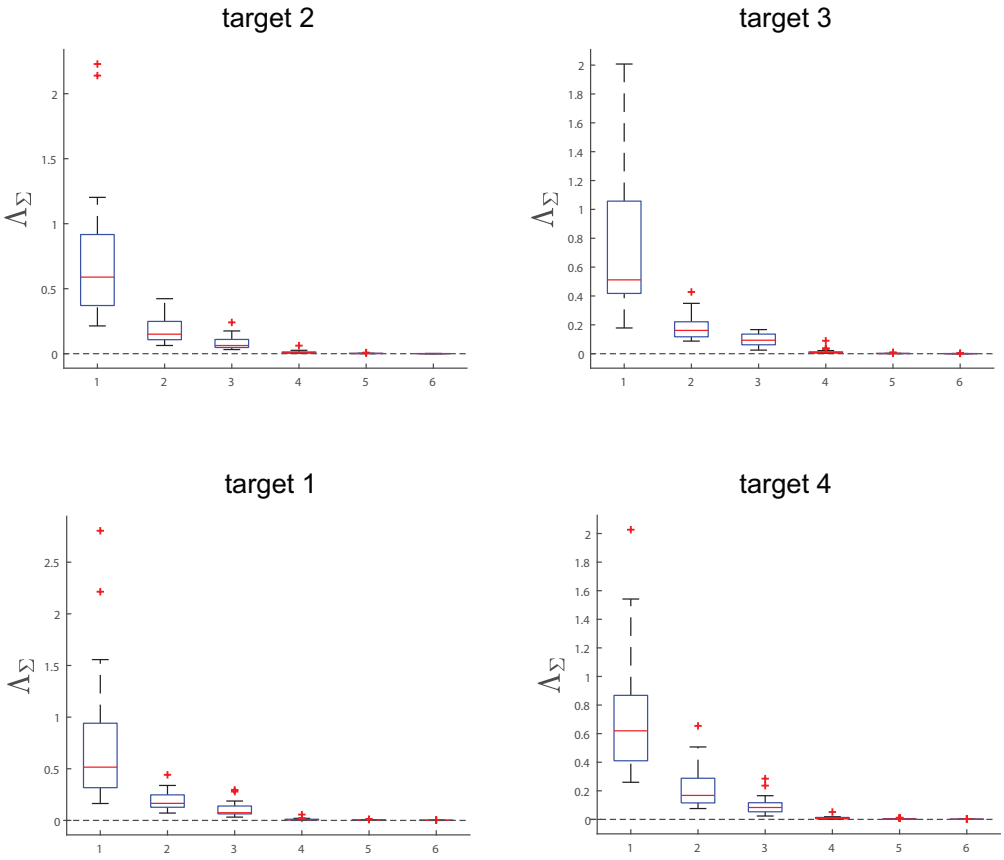


**Figure 10.** Distributions of the eigenvalues of the covariance matrix across participants for each target. The first three principal components can explain 95% of the total variance.

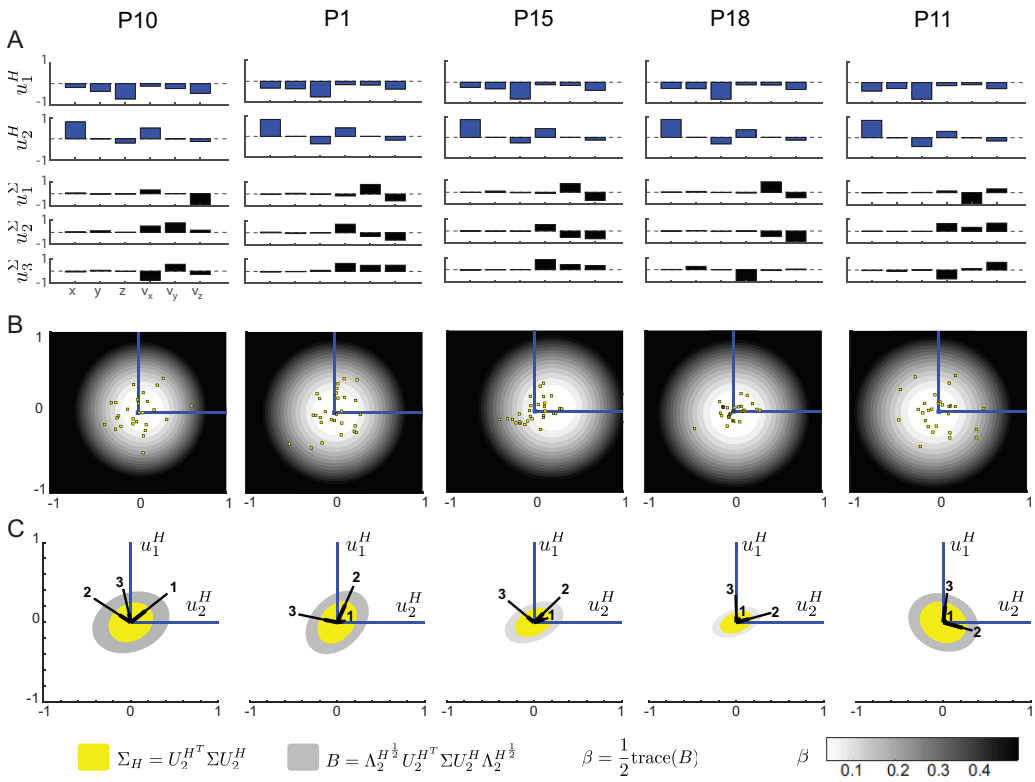


**Figure 11.** Low-dimensional representation of the *beta* index. A: principal curvature directions (blue bars) and principal variability directions (black bars). B: principal curvature plane, local action score and score-relevant perturbations (yellow square). C: principal variability directions and principal components (thin and thick black lines respectively) projected on the principal curvature plane. The three principal components contribute to the score-relevant variability which is represented with the one-standard deviation yellow covariance ellipse. The gray ellipses are a visual representation of the $\beta$ index and of the effect of the local sensitivity in amplifying the score-relevant variability (B matrix). Because $\beta$ 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.

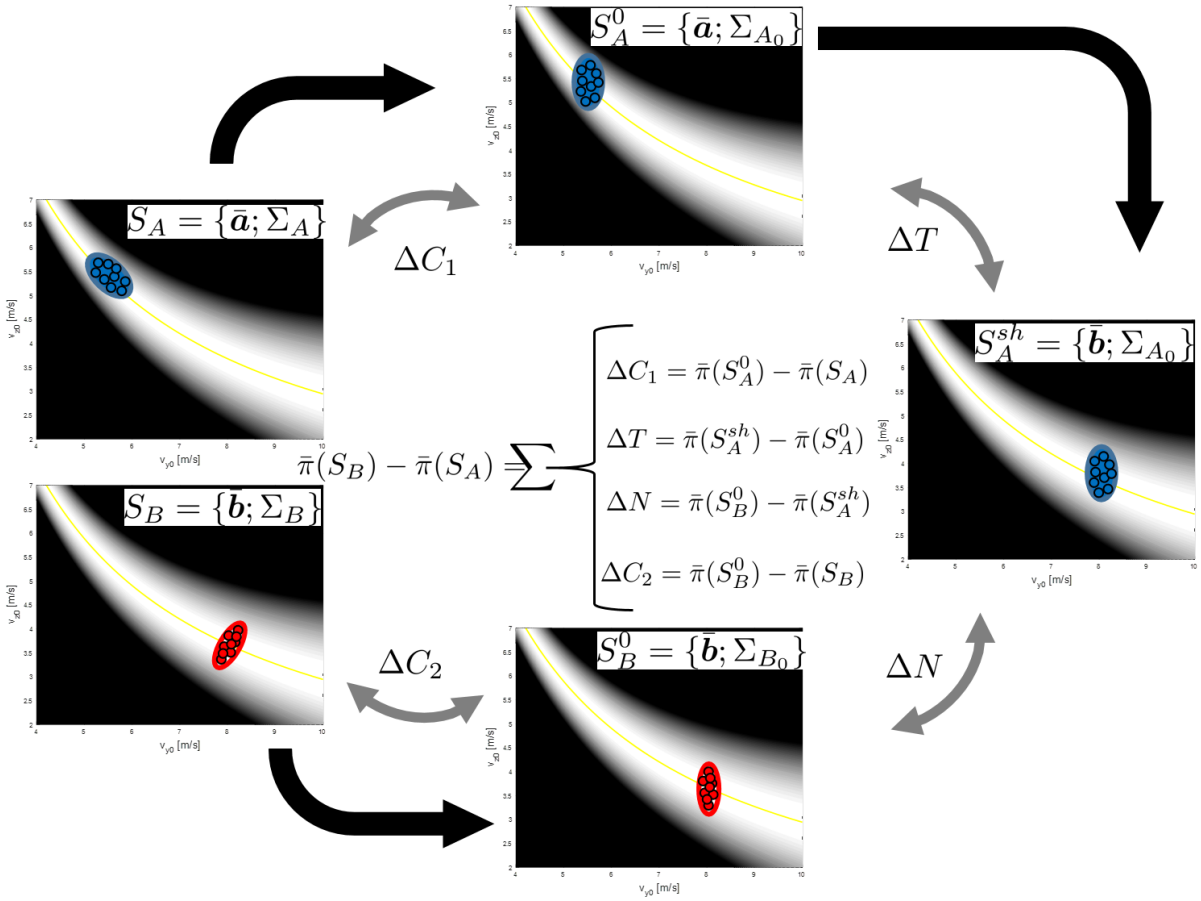


**Figure 12.** The TNC method proposed in [21]. Given two datasets, $A$ and $B$, for instance the release strategies of two different partecipants, the method requires the generation of surrogate data-sets $S^0$ (covariation-free) and $S^{sh}$ (covariation-free but shifted mean). The difference in mean score between experimental and surrogate datasets are then used to calculate the relative tolerance $\Delta T$, noise $\Delta N$ and covariation $\Delta C_1$, $\Delta C_2$ between the two strategies.