Baltic J. Modern Computing, Vol. 14 (2026), No. 2, pp. 445–461
https://doi.org/10.22364/bjmc.2026.14.2.09

Data-Driven Methods for Analyzing Non-Gaussian Variability in Human Postural Sway

Gennady CHUIKO, Yevhen DARNAPUK, Olga YAREMCHUK

Petro Mohyla Black Sea National University, Mykolaiv, Ukraine

`henadiy.chuyko@chmnu.edu.ua`, `yevhen.darnapuk@chmnu.edu.ua`,
`olga.yaremchuk.77@ukr.net`

ORCID 0000-0001-5590-9404, ORCID 0000-0002-7099-5344, ORCID 0000-0002-0891-4216

Abstract. Postural sway during quiet standing is often analyzed using diffusion- or Brownian-type models that implicitly assume near-Gaussian statistics. In practice, center-of-pressure (COP) trajectories can exhibit heavy tails, asymmetry, and intermittent excursions that may bias variance-only variability descriptors. This study presents a reproducible Python-based analysis pipeline for assessing non-Gaussian structure in COP increment series by rotating increment pairs into a PCA-decorrelated coordinate system. Using the Human Balance Evaluation Database (HBEDB) (1930 trials from 163 subjects; 60 s at 100 Hz under four standing conditions: eyes open/closed and rigid/unstable surface; three repetitions each), we compare empirical increment distributions with mean- and covariance-matched Gaussian surrogates and quantify departures using complementary diagnostics, including probability–probability comparisons, robust spread descriptors (IQR and range), and uni-/bivariate kernel density estimation. We further introduce a compact central ellipse fraction (CEF) descriptor that measures central concentration in the PCA plane and relates distributional findings to linear and nonlinear variability measures, including Poincaré-type descriptors and recurrence indices. Across the dataset, empirical increments deviate systematically from Gaussian surrogates: P–P plots show tail departures, and the 2-D density in the PCA plane is more centrally concentrated than the matched Gaussian model (CEF higher by approximately 0.08–0.10), consistent with leptokurtic (peaked and heavy-tailed) behavior. These results motivate distribution-aware preprocessing and descriptor selection for computer-based posturography and fall-risk assessment algorithms beyond variance-only summaries.

Keywords: Biomedical signal processing, Postural sway analysis, Digital signal processing, Non-Gaussian variability, Recurrence plot analysis, Computer-based balance monitoring

1 Introduction

Postural balance requires the coordinated efforts of the human body's visual, vestibular, and proprioceptive systems (Prieto et al., 1993). The proprioceptive system provides rapid muscular and joint responses to external forces. Such coherent action is proper for any posture, e.g., in the standing position. The center of gravity of a human body and the center of pressure (COP) under the feet move about a laboratory coordinate system (Collins and De Luca, 1993). This motion is commonly referred to as postural sway (Prieto et al., 1993; Collins and De Luca, 1993). As a rule, these sways are depicted via the COP trajectory on a support plane.

Meanwhile, fall-related mortality shows a pronounced peak among older adults (defined here as people aged 65 years and older; see *National Safety Council, Home and Community Overview, Injury Facts* (2023) for the latest available data from the USA). On a global scale, unintentional falls are responsible for more than 660,000 deaths each year (*World Health Organization. Falls*, 2021). As the populations of many developed countries continue to age, the burden of fatal falls is expected to grow. This raises an important question: can this trend be slowed or even reversed?

Appropriate training for older adults might be an effective prevention tool. Authors Rogers et al. (2001) reported decreases in mediolateral sway amplitude (about 9%) and mean instantaneous speed (about 13%) following a training course for older adults. Postural sways have also been used to investigate age-related changes and neurologic diseases (Prieto et al., 1993), the impact of a blood stroke (Treger et al., 2020), joint laxity (Aydın, 2017), active and passive changes in posture (Cit, 2003), during burst behavior (Picoli et al., 2019), and so forth.

The random-walk approach, also known as the Brownian motion model, is a common starting point for analyzing COP trajectory records (also called stabilograms) (Treger et al., 2020). Within this framework, stabilogram diffusion analysis (SDA) is a natural and widely used method for describing postural sway (Prieto et al., 1993; Collins and De Luca, 1993; Treger et al., 2020; Delignières et al., 2003). Despite the criticisms discussed in Delignières et al. (2003), SDA remains a valid and proper technique (Treger et al., 2020; Bao et al., 2023).

Fractional Brownian motion is a generalization of the SDA model in which the mean-square displacement grows as a power law in time. In this case, the variance of displacement depends not only on the elapsed time but also on its exponent: it increases superlinearly for persistent series and sublinearly for anti-persistent ones Delignières et al. (2003); Mandelbrot and Van Ness (1968). More generally, fractional Brownian motion treats postural sways as a kind of “i/f-type noise”, intermediate between white noise and classical Brownian motion.

However, within both models, the analysis typically emphasizes second-order (variance-based) structure, so the extent to which postural sway exhibits non-Gaussian displacement (increment) statistics is not always assessed explicitly. Meanwhile, the non-Gaussian nature of the probability distributions of many medical signals is well known. This remark also holds for other statistical methods, such as Rescaled Range Analysis (R/S) and Detrended Fluctuation Analysis (DFA), which were analyzed in (Delignières et al., 2003). One must at least know how robust the statistical parameters are.

The second point we draw attention to is the application of machine-learning and principal-component-based methods to postural sway Andò et al. (2025); Gorban and Zinovyev (2010). Rather than assuming a fixed laboratory AP–ML orientation, PCA provides a data-driven rotation of the COP increment plane into orthogonal principal axes, yielding decorrelated coordinates. This step reduces the influence of subjective and unavoidable inaccuracies in the orientation of the laboratory coordinate system. Selecting decorrelated (approximately independent) coordinates enables us to estimate two-dimensional probability density distributions of COP increments consistently across subjects and conditions. Variability is a fundamental feature of all medical signals and postural sways. To quantify this property is one of our objectives in this report. This point turned out to be close enough to the above ones.

In modern balance-assessment systems, force-platform measurements are acquired and processed by dedicated computer hardware and software. The reliability of such systems depends critically on the robustness of their data-processing algorithms. In this work, we focus on this algorithmic layer, proposing and evaluating numerical descriptors and processing steps that can be integrated into computer-based posturography and fall-risk assessment tools.

We propose a distribution-aware posturography workflow that rotates center-of-pressure (COP) increment pairs into a principal component analysis (PCA)-decorrelated coordinate system. This approach minimizes axis-orientation bias and facilitates consistent two-dimensional distribution analysis across various subjects and conditions. To achieve this, we compare the empirical statistics of COP increments with Gaussian surrogates matched for mean and covariance. This analysis incorporates complementary diagnostics, such as P–P comparisons, robust spread descriptors like the interquartile range (IQR) and range, and univariate and bivariate kernel density estimates. Furthermore, we connect the observed distribution shapes to both linear (dispersion) and nonlinear (Poincaré/recurrence) variability measures. Given that the HBEDB dataset includes repeated trials per subject across different standing conditions, we aggregate data at the subject level to ensure that individual trials are not erroneously treated as statistically independent.

This study makes three contributions to quantitative posturography. First, we provide a reproducible pipeline to test whether Gaussian assumptions are adequate for COP increment statistics by benchmarking empirical data against matched Gaussian surrogates and quantifying deviations using distributional diagnostics. Second, we introduce and operationalize bivariate distribution-shape descriptors in the PCA plane – most notably a central concentration metric based on the fraction of samples within a reference ellipse (central ellipse fraction, CEF) – to capture differences in density geometry that are not explained by variance alone. Third, using a public balance dataset spanning younger and older adults, we assess how classical dispersion descriptors, Poincaré-type measures, and recurrence-based indices relate to distribution-shape measures, highlighting which relationships remain stable and which exhibit age-linked coupling patterns.

2 Materials and Methodology

2.1 Data source and age-stratified subgroups

We used the Human Balance Evaluation Database (HBEDB) available on PhysioNet (Santos and Duarte, 2016b,a; Goldberger et al., 2000). HBEDB contains force-platform recordings of quiet standing acquired under four experimental conditions defined by vision and support surface (eyes open/closed and rigid/unstable), with three trials per condition presented in randomized order. Overall, the database provides 1930 trials from 163 participants. Each trial is 60 s long, sampled at 100 Hz, and low-pass filtered at 10 Hz. In this study, we analyzed the full dataset (all available trials).

For age-stratified analysis, participants were divided into a younger subgroup (< 60 years) and an older subgroup (≥ 60 years).

The nominal complete design would contain $163 \times 4 \times 3 = 1956$ recordings. In the public HBEDB release used in this study, 1930 trial records were available for analysis. Thus, the difference of 26 records reflects unavailable recordings in the public dataset rather than exclusions introduced by our preprocessing pipeline. We therefore analyzed all available trials.

The threshold of 60 years was used as an exploratory age-stratification choice to obtain two interpretable subgroups with sufficient sample sizes. We did not optimize this threshold or scan multiple cut-points, because such a procedure would introduce post-hoc threshold selection and increase the risk of overinterpretation. A more detailed treatment of age as a continuous variable, or the use of several age bands, should be combined with subject-level aggregation or mixed-effects modelling in future work.

In the original dataset (Santos and Duarte, 2016b), COP coordinates are provided as AP(X) and ML(Y) positions in the laboratory coordinate system. Because COP trajectories are temporally autocorrelated, the nominal sample size (6000 points) substantially overstates the number of independent observations; accordingly, statistical significance thresholds for small correlation coefficients should be interpreted with caution. Nevertheless, the observed $|corr(X, Y)|$ values were sufficiently large to indicate meaningful cross-axis coupling in the laboratory frame. A plausible reason is a slight rotation between the laboratory axes and the principal directions of sway. Therefore, following Gorban and Zinovyev (2010), we refined the data using principal component analysis (PCA) to obtain a rotated coordinate system in which the two sway components are decorrelated.

2.2 Methods

We applied principal component analysis (PCA) to orthogonalize the original COP coordinate plane, yielding decorrelated components and reducing the dependence between the anterior–posterior (AP) and mediolateral (ML) laboratory axes. In this two-dimensional setting, PCA:

- yields orthogonal principal axes,
- orders the axes by explained variance (PC1 captures maximal variance),
- diagonalizes the covariance structure, providing decorrelated coordinates that reduce sensitivity to small axis-orientation errors (Gorban and Zinovyev, 2010).

This improves the interpretability and consistency of subsequent two-dimensional distribution analysis across subjects and conditions. In the figures, the two PCA-rotated components are sometimes labelled AP and ML for visual continuity; however, they should be interpreted as rotated principal components rather than as the original anatomical/laboratory AP and ML axes.

After PCA rotation, we analyzed displacements between successive positions (increment series), which are commonly used in stabilogram diffusion analysis (SDA) and related approaches. For each trial and component, we examined increment distributions using the Shapiro–Wilk normality test (Shapiro and Wilk, 1965) and graphical diagnostics (box plots and probability–probability (P–P) plots) (Field, 2024). Because the increment sequences are autocorrelated and contain $n = 5999$ successive increments per trial, Shapiro–Wilk p-values were interpreted as sensitivity indicators rather than as i.i.d. assumption checks. Nonparametric probability density functions were estimated using kernel density estimation (KDE). In the reported implementation, an Epanechnikov kernel was used for the KDE-based visualizations, with bandwidths selected by a normal-reference rule-of-thumb (Silverman-type rule, coefficient $k = 1.06$) (Zielinski et al., 2018; Silverman, 2018; Moraes et al., 2021). The term “normal-reference” refers to the bandwidth-selection heuristic only and does not imply that the empirical COP increment distributions are assumed to be Gaussian. In addition to univariate densities, we estimated two-dimensional probability density functions (2D-PDFs) for the joint increment distribution in the PCA plane.

To quantify departures from a Gaussian reference model in the 2D case, we introduced a compact distribution-shape descriptor based on the one-sigma region: the central ellipse fraction (CEF), defined as the fraction of samples whose Mahalanobis distance in the PCA plane satisfies $d_M \leq 1$, where the covariance structure of the corresponding Gaussian model determines the reference ellipse. This measure captures differences in central concentration between empirical 2D distributions and their mean- and covariance-matched Gaussian surrogates.

A complementary analysis focused on variability descriptors. We computed three dispersion descriptors (standard deviation, interquartile range, and range) (Field, 2024) and three additional descriptors: recurrence rate (RR) from recurrence analysis (Marwan et al., 2018) and Poincaré-plot measures SD1 and SD2 (Golińska, 2013). We assessed cross-correlations among the six descriptors (including correlation magnitudes) and compared descriptor distributions and density estimates between components (AP vs ML) and between younger and older cohorts. All analyses and figure generation were implemented in Python using standard scientific computing libraries (Hudson et al., 2022).

3 Results

3.1 Deviations of the factual distributions from the standard Gaussian reference

To begin, we transformed the COP position series (X, Y)(Santos and Duarte, 2016c) into an increment (displacement) series between successive samples, $(\delta x, \delta y)$. In the laboratory coordinate system, the increment components exhibited cross-axis coupling (a nonzero Pearson correlation), consistent with a slight misalignment between the laboratory axes and the dominant sway directions. Therefore, we applied principal component analysis (PCA) to the increment pairs and rotated the laboratory frame into a PCA frame in which the two components are linearly decorrelated.

In the PCA frame, the components have near-zero linear correlation. However, decorrelation does not generally imply statistical independence, particularly when the variability is non-Gaussian. Accordingly, we treat the PCA components as linearly decorrelated variables while allowing higher-order dependencies, and we estimate joint distributions directly rather than factorizing them into products of marginals. Because PCA orders axes by explained variance, this rotation also yields a variance-maximizing coordinate system that improves the interpretability of subsequent two-dimensional distribution analysis.

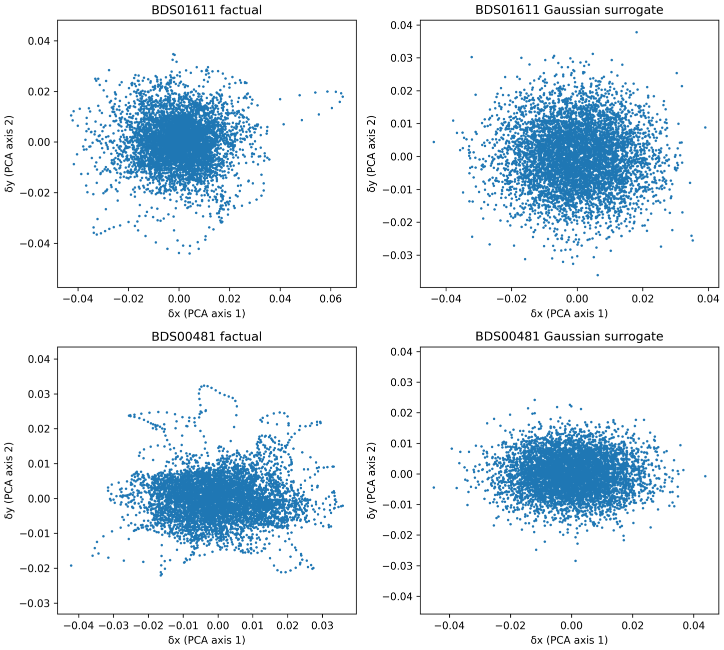

Fig. 1. Scatter plots of COP increment pairs (δx, δy) in the PCA-rotated coordinate system (axes decorrelated). Two example trials are shown: BDS01611 (27 years; upper row) and BDS00481 (69 years; lower row). Empirical series are shown on the left; the right panels show Gaussian surrogate series generated to match the empirical mean vector and covariance matrix while preserving trial length, enabling visual comparison of tail behavior and anisotropy. Axis limits are scaled separately across panels to preserve local point-cloud visibility; therefore, the figure should be interpreted qualitatively in terms of distributional shape, tail organization, and empirical-versus-surrogate differences rather than absolute plotted area.

Figure 1 shows scatter plots of successive COP increments in the PCA-rotated plane (PC1–PC2). For comparison, we generated Gaussian surrogate samples from a bivariate normal distribution with the same empirical mean vector and covariance matrix and the same number of increments. The empirical and surrogate point clouds are broadly similar in overall scale and anisotropy; however, the empirical extremes often form structured arcs/clusters consistent with temporal dependence, whereas surrogate extremes are more spatially diffuse.

To quantify differences in central concentration, we define the central ellipse fraction (CEF) as the proportion of samples lying inside the $1 - \sigma$ ellipse in the PCA plane, equivalently, those with Mahalanobis distance $d_M \leq 1$ under the matched Gaussian reference. Shapiro–Wilk tests rejected normality for all AP trials (1930/1930) and most ML trials (1926/1930) at $\alpha = 0.05$. Because increments inherit autocorrelation from COP trajectories and $n \approx 6000$ per trial yields very high test power, we treat Shapiro–Wilk as a supporting diagnostic alongside P–P plots and KDE-based density estimates rather than as a standalone proof of non-Gaussianity.

As illustrated in Figure 2, probability–probability (P–P) plots compare the empirical cumulative distribution with that of a reference Gaussian distribution (Field, 2024). In our data, the curves show systematic departures from the diagonal (the Gaussian reference), most visibly toward the tails of the distribution, consistent with deviations in kurtosis and, in some trials, mild asymmetry. Together with the Shapiro–Wilk test results, these plots support the conclusion that the increment/displacement distributions are not exactly normal.

The two examples in Figure 2 also show that the degree of deviation is not identical across trials or components. In trial BDS01684, both components show visible systematic departures from the diagonal, especially away from the center of the distribution, whereas in trial BDS00995 the ML component remains much closer to the Gaussian reference over most of the probability range. This trial-to-trial heterogeneity is one reason why we rely on several complementary diagnostics rather than on a single normality test or a single visual summary.

A second useful view of the same phenomenon is provided by the box-and-whisker representation in Figure 3. Whereas the P–P plots emphasize distributional agreement with a Gaussian reference over the full probability range, the box plots separate the central spread from the most extreme values. In both illustrated trials, the AP component occupies a wider range than the ML component, and the most pronounced differences are concentrated in the tails rather than in the median location.

Note that the empirical distributions can look close to a Gaussian model for small displacements near the center; however, differences become more apparent for larger displacements (extremes). These extremes are visible in the left-hand scatter plots of Figure 1. They are reflected in the box-and-whisker plots in Figure 3. Also note that the most extreme values are often not perfectly balanced around the median (and thus around zero), indicating that tail behavior can be asymmetric even when the central bulk is close to symmetric. The ranges of AP sways exceed those of ML sways in the shown examples (Figure 3). The presence of extreme values can strongly inflate the range and standard deviation and may also broaden the IQR (Field, 2024). We will return to this issue when we consider the variability descriptors of postural sway.

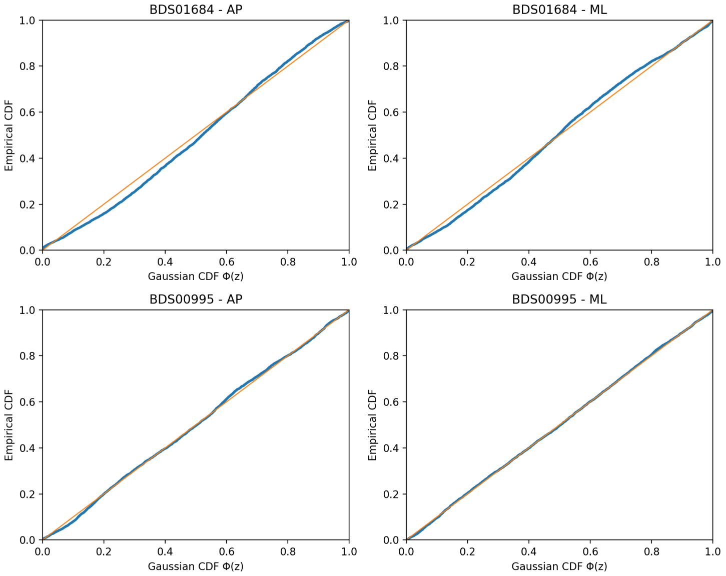

Fig. 2. Probability–probability (P–P) plots comparing the empirical CDF of COP increment samples with the CDF of a Gaussian reference. AP is shown on the left and ML on the right. The upper row shows trial BDS01684 (younger participant), and the lower row shows trial BDS00995 (older participant). The thin diagonal line is the identity line $y = x$, indicating perfect agreement with the Gaussian reference; systematic departures from this line indicate non-Gaussianity (especially in the tails).

The two examples also illustrate that extreme values need not occur symmetrically in the positive and negative directions. In BDS01684, the AP component shows a pronounced negative tail together with several large positive excursions, while the corresponding ML component remains more compact. In BDS00995, the AP component again exhibits a broader spread and more numerous outliers than the ML component. This pattern is consistent with the impression from Figure 1 that the empirical increment clouds differ from Gaussian surrogates primarily through their tail structure and central concentration rather than through large shifts of the mean.

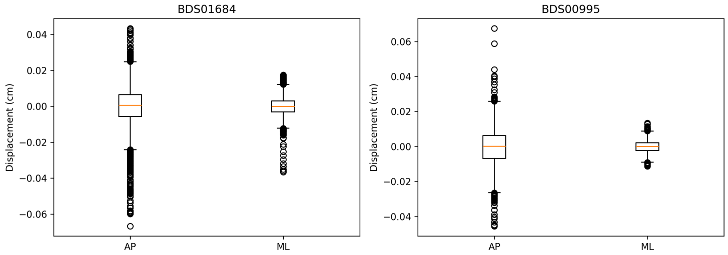

Fig. 3. Box-and-whisker plots of the two PCA-rotated COP increment components (PC1 and PC2; labeled AP and ML in the plot) for two example trials (BDS01684 and BDS00995). Boxes indicate the interquartile range (IQR) with the median shown as a horizontal line; whiskers extend to the most extreme values within $1.5 \times IQR$ (Tukey convention); points beyond the whiskers are plotted as outliers.

3.2 Kernel density estimations

Kernel density estimation (KDE) is a method for estimating the probability density function (PDF) of a distribution (Zielinski et al., 2018; Silverman, 2018). The following two choices are most important within KDE:

- the choice of appropriate kernel functions,
- the right choice of the bandwidth (also called the smoothing parameter (Zielinski et al., 2018)).

Kernel density estimation (KDE) was used as a descriptive visualization tool for empirical and surrogate distributions. The KDE itself is nonparametric; therefore, it does not require the underlying data to be normally distributed. In the reported figures, KDEs were computed with an Epanechnikov kernel and a normal-reference rule-of-thumb bandwidth. The term "normal-reference" refers only to the bandwidth-selection heuristic and does not imply that the empirical COP increment distributions are assumed to be Gaussian. The Gaussian assumption is used only for constructing the mean- and covariance-matched surrogate samples, not for estimating the empirical density.

The bandwidth was computed as

$$h = k \cdot \min(\mathrm{StD}, 0.75 \times \mathrm{IQR}) \cdot n^{-1/5},$$

where h denotes the bandwidth, StD is the sample standard deviation, IQR is the interquartile range, and n is the sample size. We used $k = 1.06$ to keep the smoothing rule stable and comparable across empirical and surrogate distributions. Kernel choice mainly affects visual smoothness, whereas bandwidth selection controls the main bias–variance trade-off in the estimated density.

To visualize and compare the distributions of the central-ellipse descriptor for factual and Gaussian surrogate data, we used violin plots, which display both the empirical density and basic summary statistics.

Figure 4 shows that the point clouds in Figure 1 differ not only in the periphery but also in the central region. Specifically, the central ellipse fraction (CEF)—the proportion of increment samples falling inside the $1 - \sigma$ ellipse in the PCA plane—is

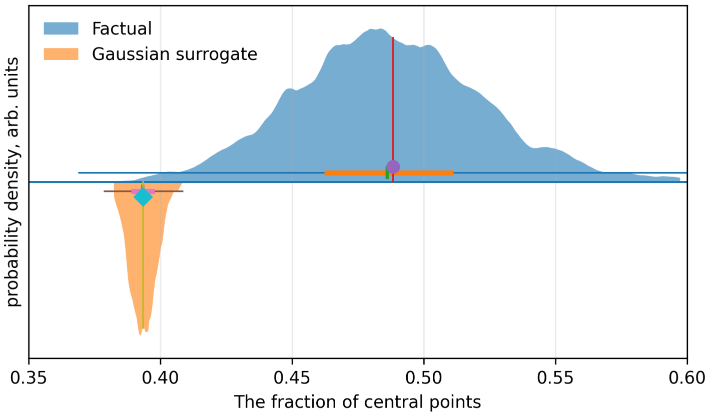

Fig. 4. Split-violin distributions of the central ellipse fraction (CEF) across trials, comparing empirical COP increment clouds (upper half, blue) with mean- and covariance-matched Gaussian surrogates (lower half, orange). The violin shapes were estimated using Epanechnikov KDE with a rule-of-thumb bandwidth. CEF is defined as the fraction of increment pairs that fall inside the $1-\sigma$ reference ellipse in the PCA plane (equivalently, Mahalanobis distance $d_M \leq 1$ under the matched Gaussian model). Points mark the mean CEF, and vertical lines mark the median CEF for each distribution.

higher for the empirical data than for the mean- and covariance-matched Gaussian surrogates, indicating stronger central concentration in the real distributions. The violin shapes in Figure 4 are kernel density estimates of the CEF distributions across trials for the factual and surrogate series (Zielinski et al., 2018; Silverman, 2018; Moraes et al., 2021). Because CEF is computed from the bivariate increment vectors after rotation to the PCA frame, it reflects the joint distribution in the PCA plane; however, the KDE shown in Figure 4 is applied to the univariate descriptor values rather than to the full two-dimensional increment clouds.

As a bandwidth-sensitivity check for Figure 4, we repeated the CEF density estimation using leave-one-out likelihood cross-validated bandwidths. For the factual CEF distribution, the rule-of-thumb bandwidth was $h = 0.00822$, whereas the cross-validated bandwidth was $h = 0.02236$; for the Gaussian-surrogate CEF distribution, the corresponding values were $h = 0.00116$ and $h = 0.00207$. Thus, cross-validation selected larger bandwidths and therefore smoother density estimates in both cases. Importantly, the qualitative separation between the factual and Gaussian-surrogate CEF distributions remained unchanged: empirical CEF values stayed shifted toward higher central concentration. Therefore, the rule-of-thumb bandwidth was retained in the main figures to keep the smoothing rule stable and comparable across all trial-level distributions.

Figure 5 highlights qualitative differences between the empirical distributions and the mean- and covariance-matched Gaussian surrogate distributions in the PCA increment plane. The surrogate densities retain the expected smooth, approximately elliptical contour structure, whereas the empirical KDE contours can exhibit sharper central con-

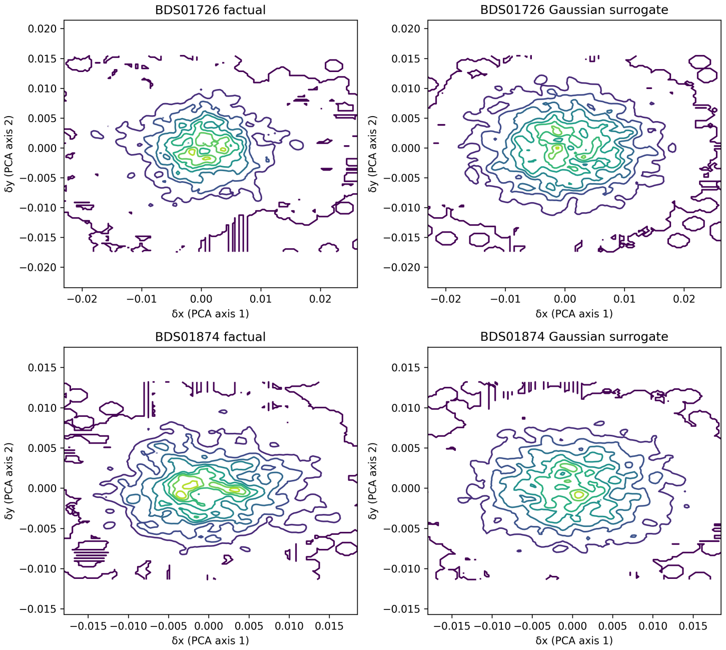

Fig. 5. Contour maps of the bivariate KDE of COP increment pairs in the PCA-rotated increment plane (decorrelated axes). Two example trials are shown: BDS01726 (31 years; upper row) and BDS01874 (62 years; lower row). Empirical densities are shown in the left panels; the right panels show KDEs of Gaussian surrogate samples generated to match the empirical mean vector and covariance matrix (same trial length), enabling visual comparison of central concentration, tail structure, and anisotropy. Axis limits are scaled separately across panels to preserve contour readability; therefore, the figure should be interpreted qualitatively in terms of density shape, central concentration, and contour regularity rather than absolute spatial extent.

centration and increasingly irregular deviations from ellipses toward the periphery. In some trials, the empirical density is close to unimodal. Still, in others, it shows noticeable asymmetry and/or secondary features (e.g., the empirical BDS01874 panel), consistent with non-Gaussian shape effects beyond second-order covariance structure. These observations are consistent with the central-ellipse-fraction results, which indicate that the empirical increments exhibit systematically stronger central concentration than the Gaussian surrogates.

3.3 Variability and its descriptors

Linear statistical descriptors of variability include the standard deviation (SD), interquartile range (IQR), and range (R) (Field, 2024). We also computed two Poincaré-plot descriptors (SD1 and SD2) (Golińska, 2013) and the recurrence rate (RR) from recurrence analysis (Marwan et al., 2018). Figure 6 illustrates these representations for a typical PCA-rotated COP increment series. The elongated Poincaré clouds indicate SD2 $\gg$

SD1, i.e., variability along the identity line (longer-term component) dominates the short-term component orthogonal to it. A small number of extreme points at the plot ends can inflate SD2 (and, to a lesser extent, SD1).

Recurrence plots were constructed as 1D amplitude recurrence matrices using the criterion $\|x_i - x_j\| \le \epsilon$. In our implementation, ϵ was set to the normal-reference (Silverman) bandwidth (Eq. (1), $k = 1.06$) computed for each series, and RR summarizes the fraction of recurrent pairs. Because RR aggregates over all index pairs, it is typically less sensitive to a small number of outliers than range-based dispersion measures. However, extreme points can still affect RR through their distances to other samples.

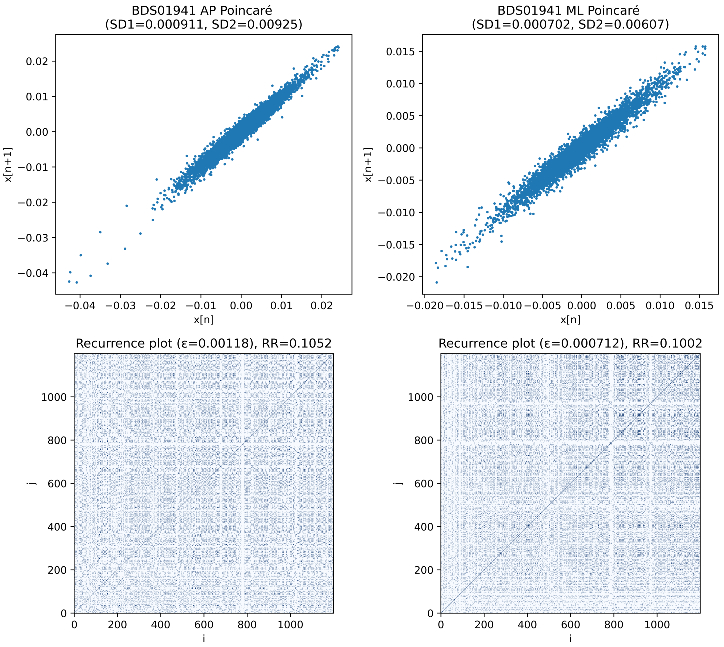

Fig. 6. Example Poincaré plots (top row) and recurrence plots (bottom row) for the PCA-rotated COP increment series from HBEDB record BDS01941 (23 years). Left panels: first component (labeled AP); right panels: second component (labeled ML). Poincaré plots show $x[n]$ versus $x[n+1]$, with SD1 and SD2 reported in the panel titles. Recurrence plots visualize the binary recurrence matrix defined by $\|x_i - x_j\| \le \epsilon$, where ϵ is set to the normal-reference (Silverman) bandwidth ($k = 1.06$) computed for each series; the corresponding recurrence rate (RR) is shown. For visualization, recurrence matrices are downsampled to 1200 points, whereas RR is computed on the full series. For the Poincaré plots, axis limits are scaled separately for the components labeled AP and ML to preserve the visible geometry of each component. The recurrence plots use the same index scale and are compared by recurrence structure and recurrence rate.

We summarized the distributions of three variability descriptors (SD2, StD, and RR) using split violin plots (Figure 7). Across both age groups, SD2 and StD are consistently higher in the AP direction than in the ML direction. The SD2 and StD distributions are very similar, consistent with their near-perfect positive association, whereas RR shows a weaker and predominantly inverse association with these dispersion measures. In the older subgroup, the distributions of SD2 and StD appear broader, with more pronounced tails, suggesting greater heterogeneity in variability across trials.

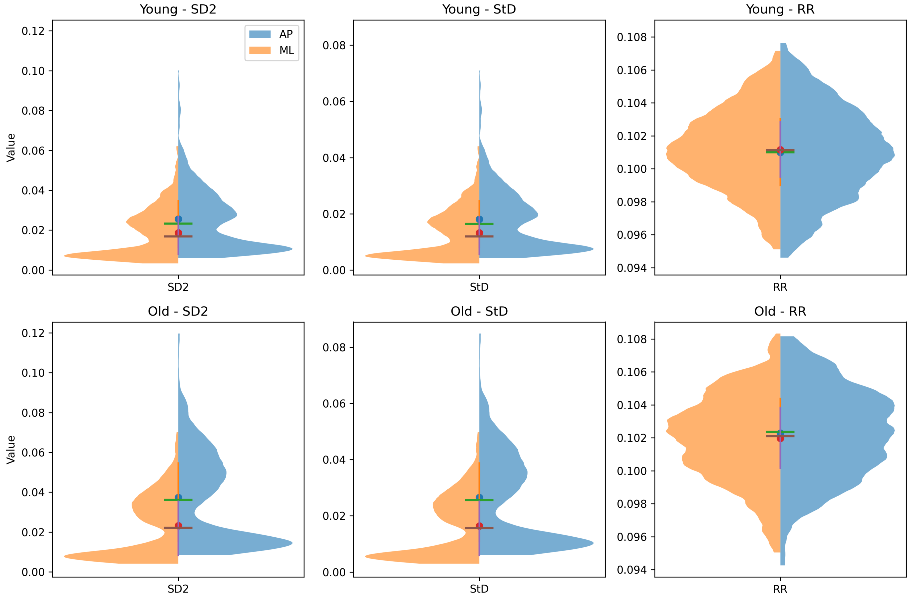

Fig. 7. Split violin plots (KDE-based) for three variability descriptors: SD2 (left column), StD (middle column), and RR (right column). The upper row shows the younger group; the lower row shows the older group. Each violin is split by direction: AP (right half) and ML (left half). Points and short line markers indicate the mean, median, and interquartile range for each half.

We examined linear associations among six variability descriptors (SD2, SD1, StD, IQR, Range, and RR) using Pearson correlation coefficients computed across trials within each age-stratified subgroup and direction. To highlight the descriptor that exhibited the clearest group contrast, we summarize in Figure 8 the correlations between the short-term variability measure SD1 and each of the remaining descriptors (SD2, StD, IQR, Range, and RR), separately for AP and ML sways in the younger and older groups. For compact visualization, Figure 8 reports absolute correlation magnitudes ($|r| \in [0, 1]$); therefore, inverse relationships (e.g., negative SD1–RR correlations) appear as large $|r|$ values, but their signs are discussed in the text rather than encoded in the bars.

Across most descriptor pairs, the observed correlations were consistently high and showed little dependence on direction. In contrast, SD1 showed weaker coupling with the other descriptors in the younger subgroup than in the older subgroup, suggesting

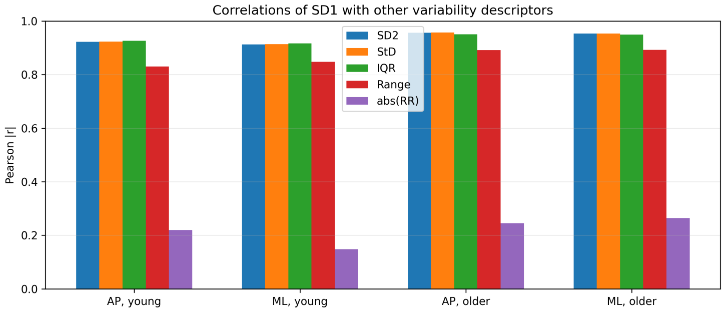

Fig. 8. Absolute Pearson correlations ($|r|$) between SD1 and SD2, StD, IQR, Range, and RR for AP/ML sway in younger and older groups. SD1 shows strong coupling with dispersion descriptors and weaker coupling with RR; age-related differences are shown descriptively.

that short-term variability may be less redundant with long-term and distributive variability measures in younger participants. Because each coefficient is estimated from a finite number of trials per subgroup, we treat these correlation patterns as effect-size summaries and interpret subgroup differences as exploratory unless confirmed by confidence intervals and an explicitly multiple-comparison-controlled inference procedure.

4 Discussion and limitations

The observed deviations from Gaussian displacement statistics indicate that COP quiet-standing trajectories can contain intermittent excursions that are not well captured by variance-only descriptors. From an algorithmic perspective, this matters because several widely used posturographic measures (including diffusion-based scaling and standard dispersion/correlation summaries) can be sensitive to heavy tails and to a small proportion of extreme samples. The surrogate comparison provides a pragmatic baseline: if the Gaussian surrogate preserves the gross scale of motion yet fails to reproduce central concentration or tail behavior, then distribution shape – not only variance – contributes materially to the signal structure.

The central-ellipse fraction descriptor complements tail-focused criteria by quantifying whether sway is centrally concentrated (a sharp peak) or more broadly distributed and structurally non-elliptical in the PCA plane. In practice, this can distinguish trajectories with similar RMS dispersion but different density geometry, which is relevant for classifier design and robustness testing of embedded posturography software. Likewise, differences in coupling patterns among variability descriptors across age groups suggest that “equivalent” variability metrics may not remain interchangeable under aging; this motivates reporting multiple descriptors and explicitly checking correlations rather than assuming redundancy. In particular, the near-perfect SD2–StD coupling and the weaker inverse association with RR suggest that recurrence captures aspects of temporal structure that are not reducible to simple dispersion, reinforcing the value of reporting both dispersion and recurrence-based descriptors.

Several limitations should be acknowledged. First, although the study uses the full HBEDB corpus, the data includes repeated trials per subject across conditions; treating trials as independent can inflate apparent significance. Subject-level aggregation (e.g., within-subject medians) and mixed-effects models would strengthen inferential claims, and future work should emphasize effect sizes with confidence intervals rather than p-value-only interpretation. Second, the analysis is restricted to the HBEDB acquisition protocol (60-s trials sampled at 100 Hz and low-pass filtered at 10 Hz). Distribution shape and recurrence metrics may vary with task demands, recording duration, sampling rate, filtering, and nonstationarity, so generalization to other protocols should be made cautiously. Third, Gaussian surrogates test a specific null model (matched mean vector and covariance); alternative surrogates (e.g., phase-randomized, block-resampled, or AR-matched) may be needed to separate distributional non-Gaussianity from temporal dependence. Finally, recurrence-plot outcomes depend strongly on parameter choices, especially the distance threshold ϵ; routine reporting of parameter sensitivity would improve reproducibility.

Future work should extend the evaluation to additional datasets and stance conditions, include parameter-robust recurrence analysis, and assess whether distribution-shape descriptors improve predictive models for fall risk beyond standard dispersion measures.

5 Conclusions

Principal component analysis (PCA) provided a rotated coordinate system in which the two COP increment components were linearly decorrelated and ordered by variance (Gorban and Zinovyev, 2010). This PCA-based rotation enabled consistent two-dimensional analysis in the decorrelated plane and supported the estimation of bivariate probability density functions for empirical and mean-/covariance-matched Gaussian surrogate distributions (Figure 5).

In this report, we demonstrated the non-Gaussian nature of the actual distributions (see Figs. 1, 2, 3, 4, 5) compared to the Gaussian one using complementary diagnostics. In particular, Shapiro–Wilk testing rejected normality for all AP trials (1930/1930) and for almost all ML trials (1926/1930) at $\alpha = 0.05$, while P–P plots and KDE-based estimates show systematic departures concentrated toward the tails and in the density shape. Non-Gaussianity manifests as skewness, excessive kurtosis, and, primarily, systematic outliers in empirical distributions. The fraction of outliers appears to be relatively small, ranging from 3.5% to 11%. However, one should also account for their sufficiently large deviations from the mean values there.

Nevertheless, the distributions are mostly unimodal, roughly symmetric around the mean, and have similar means, medians, and modes, so the central bulk can appear close to Gaussian even when the overall distribution is not. Therefore, while non-Gaussianity is statistically pervasive across trials, it is largely expressed through tail behavior and central-concentration differences relative to a mean/covariance-matched Gaussian model. In this sense, variance-based modelling assumptions (including the Fractional Brownian motion approach) remain useful for describing the bulk scaling-dispersion structure, but tail-sensitive analyses and robustness claims should explicitly account for heavy-tailed excursions and distribution shape.

Six descriptors of variability (three linear and three non-linear) are strongly correlated with one another. Subtle age-related hallmarks were observed in some of them (see Figs. 7 and 8). A simple central-occupancy descriptor for bivariate displacement distributions was introduced and compared with standard nonlinear measures, such as the Recurrence Ratio, providing a compact way to quantify distribution-shape differences that are not captured by variance-only summaries.

References

Andò, B., Baglio, S., Manenti, M., Finocchiaro, V., Marletta, V., Rajan, S., Nehary, E. A., Dibilio, V., Zappia, M., Mostile, G. (2025). Investigating performance of an embedded machine learning solution for classifying postural behaviors, *Sensors* **25**, 4262.

Aydın, E. (2017). Postural balance control in women with generalized joint laxity, *Turkish Journal of Physical Medicine and Rehabilitation* **63**, 259–265.

Bao, W., Tan, Y., Yang, Y., Chen, K., Liu, J. (2023). Correlation of balance posturographic parameters during quiet standing with the berg balance scale in patients with parkinson's disease, *BMC Neurology* **23**.

Cit (2003). *Circulatory response to passive and active changes in posture*, Pennsylvania State University.

Collins, J. J., De Luca, C. J. (1993). Open-loop and closed-loop control of posture: A random-walk analysis of center-of-pressure trajectories, *Experimental Brain Research* **95**, 308–318.

Delignlères, D., Deschamps, T., Legros, A., Caillou, N. (2003). A methodological note on non-linear time series analysis: Is the open-and closed-loop model of collins and de luca (1993) a statistical artifact?, *Journal of Motor Behavior* **35**, 86–96.

Field, A. (2024). *Discovering Statistics Using IBM SPSS Statistics*, 6 edn, Sage Publications.

Goldberger, A. L., Amaral, L. A. N., Glass, L., Hausdorff, J. M., Ivanov, P. C., Mark, R. G., Mietus, J. E., Moody, G. B., Peng, C.-K., Stanley, H. E. (2000). Physiobank, physiotoolkit, and physionet: Components of a new research resource for complex physiologic signals, *Circulation* **101**(23), e215–e220.

Golińska, A. K. (2013). Poincaré plots in analysis of selected biomedical signals, *Studies in Logic, Grammar and Rhetoric* **35**, 117–127.

Gorban, A. N., Zinovyev, A. Y. (2010). *Principal Graphs and Manifolds*, Handbook of Research on Machine Learning Applications and Trends, pp. 28–59.

Hudson, D., Wiltshire, T. J., Atzmueller, M. (2022). multisyncpy: A python package for assessing multivariate coordination dynamics, *Behavior Research Methods* **55**, 932–962.

Mandelbrot, B. B., Van Ness, J. W. (1968). Fractional brownian motions, fractional noises and applications, *SIAM Review* **10**, 422–437.

Marwan, N., Webber, C. L., , E., Viana, R. L. (2018). Introduction to focus issue: Recurrence quantification analysis for understanding complex systems, *Chaos An Interdisciplinary Journal of Nonlinear Science* **28**.

Moraes, C. P., Fantinato, D. G., Neves, A. (2021). Epanechnikov kernel for pdf estimation applied to equalization and blind source separation, *Signal Processing* **189**, 108251.

National Safety Council, Home and Community Overview, Injury Facts (2023). `https://injuryfacts.nsc.org/home-and-community/home-and-community-overview/introduction/`

Picoli, S., Santos, E. S. D., Deprá, P. P., Mendes, R. S. (2019). Quantifying postural sway dynamics using burstiness and interevent time distributions, *The European Physical Journal B* **92**(7).

Prieto, T., Myklebust, J., Myklebust, B. (1993). Characterization and modeling of postural steadiness in the elderly: a review, *IEEE Transactions on Rehabilitation Engineering* **1**, 26–34.

Rogers, M. E., Fernandez, J. E., Bohlken, R. M. (2001). Training to reduce postural sway and increase functional reach in the elderly, *Journal of Occupational Rehabilitation* **11**, 291–298. `https://pubmed.ncbi.nlm.nih.gov/11826729/`

Santos, D. A., Duarte, M. (2016a). Human balance evaluation database. `https://physionet.org/content/hbedb/`

Santos, D. A., Duarte, M. (2016b). A public data set of human balance evaluations, *PeerJ* **4**, e2648.

Santos, D. A., Duarte, M. (2016c). A public data set of human balance evaluations.

Shapiro, S. S., Wilk, M. B. (1965). An analysis of variance test for normality (complete samples), *Biometrika* **52**(3-4), 591–611.

Silverman, B. (2018). *Density Estimation for Statistics and Data Analysis*, Routledge.

Treger, I., Mizrachi, N., Melzer, I. (2020). Open-loop and closed-loop control of posture: Stabilogram-diffusion analysis of center-of-pressure trajectories among people with stroke, *Journal of Clinical Neuroscience* **78**, 313–316.

World Health Organization. Falls (2021). `https://www.who.int/news-room/fact-sheets/detail/falls`

Zielinski, W., Kuchar, L., Michalski, A., Kazmierczak, B. (eds) (2018). *Kernel Density Estimation and Its Application*, Vol. 23, ITM Web of Conferences.

Received February 5, 2026 , revised May 7, 2026, accepted June 27, 2026