



On Projections to Linear Subspaces

Erik Thordsen^(✉)  and Erich Schubert 

TU Dortmund University, Otto-Hahn-Straße 14, 44227 Dortmund, Germany
{erik.thordsen,erich.schubert}@tu-dortmund.de

Abstract. The merit of projecting data onto linear subspaces is well known from, e.g., dimension reduction. One key aspect of subspace projections, the maximum preservation of variance (principal component analysis), has been thoroughly researched and the effect of random linear projections on measures such as intrinsic dimensionality still is an ongoing effort. In this paper, we investigate the less explored depths of linear projections onto explicit subspaces of varying dimensionality and the expectations of variance that ensue. The result is a new family of bounds for Euclidean distances and inner products. We showcase the quality of these bounds as well as investigate the intimate relation to intrinsic dimensionality estimation.

1 Introduction

The probably most important research on linear subspace projections was written by Pearson in his 1901 paper on Principal Component Analysis (PCA). The concept of PCA explains how the variance of a data set can be decomposed into orthogonal components, each of which covers the maximum amount of variance. This fundamental result has been employed in many fields including dimensionality reduction, clustering [1], intrinsic dimensionality estimation [5], and many more. The decomposition also implies linear projections that preserve the least amount of variance. Yet, it yields little information on the less tangible middle ground of random projections. The Johnson-Lindenstrauss lemma shows that random projections can preserve distances well, and the effect of random projections on, e.g., intrinsic dimensionality [6] has also been explored in the past. But we could not find literature on the effect of random projections on the variance itself. In this paper, we investigate the effect on a projected point's squared norm which entails effects on the variance of the data set. The arising bounds for the Euclidean distance as well as for inner products are explored in Sect. 2. The projections required for these bounds rely on the normal vectors of the linear subspace on which we project, which are drawn from the data set itself. Using measures based on points from the data set to assess boundaries on norms is a concept already employed in, e.g., spatial indexing. Methods like LAESA [7] use so-called pivot/reference/prototype points and the triangle inequality to prune

Part of the work on this paper has been supported by Deutsche Forschungsgemeinschaft (DFG), project number 124020371, within the Collaborative Research Center SFB 876 “Providing Information by Resource-Constrained Analysis”, project A2.

the data set during spatial queries. Tree-based methods like the Balltree [8] use the triangle inequality to exclude entire subtrees, while permutation based indexing [3, 14] uses the relative closeness to reference points to partition the data. The central points in these approaches fulfill a role equivalent to pivots. Using pivots for random projections, however, yields fundamentally stronger pruning capabilities, as discussed in Sect. 2. In Sect. 3, we analyze the expected values of variance preserved by random projections. These expectations are closely related to PCA, yet costly to compute exactly. To compensate for the computational cost and fathom the relation to eigenvalues we propose an approximation of the expected values in terms of eigenvalues. The expected values are related to the Angle-Based Intrinsic Dimensionality (ABID) estimator [13]. We explore the relationship in Sect. 4, which leads to a tangible link between indexing complexity and intrinsic dimensionality. To highlight the practical implications as well as showcase the efficacy of the introduced bounds we propose a very simple index and our empirical results in Sect. 5. Lastly, we close with a summary of this paper and a short outlook on future research in Sect. 6.

In this paper, we denote the i -th eigenvalue of some matrix M with $\lambda_i^{(M)}$. We do not care about the specific order of eigenvalues but assume that corresponding eigenvalues of matrices that admit the same eigenvectors are in the same order. We write M^c as an abbreviation for $V\Lambda^cV^T$ where V is the matrix containing the eigenvectors of M as columns and Λ^c is the diagonal matrix containing $(\lambda_i^{(M)})^c$ on the diagonal. We write $C(X)$ for the covariance matrix of data sets X where we assume X to be origin-centered unless otherwise specified. We denote the normalizations of vectors x and data sets X with \tilde{x} and \tilde{X} , respectively. Whenever Euclidean spaces and distances are discussed, the dot product is implied by the inner product.

2 Pivotal Bounds in Euclidean Spaces

We consider linear subspace projections of *query points* onto the linear subspace spanned by (not necessarily orthogonal) *pivots* or *reference points* $\{r_1, \dots, r_k\}$, $k \leq d$ drawn from the same distribution as the analyzed data set, e.g., by choosing them from the data set itself. In the case of affine subspace projections, both the query and reference points are shifted by a *center* point c . We assume all (shifted) reference points to be linearly independent. Otherwise, we discard reference points until linear independence holds. The projection $\pi(x-c; r_1-c, \dots, r_k-c)$ of some shifted query point $x-c$ onto the affine subspace (shortened to $\pi(x-c)$ whenever the choice of reference points is clear) is then given by

$$\pi(x-c) = \sum_{i=1}^k \langle x-c, \hat{r}_i \rangle \hat{r}_i \quad (1)$$

where the \hat{r}_i are the normalized orthogonal vectors obtained from the Gram-Schmidt process applied to the r_i-c . These can be recursively computed from

$$\hat{r}_1 = \frac{r_1 - c}{\|r_1 - c\|} \quad \hat{r}_i = \frac{(r_i - c) - \sum_{j=1}^{i-1} \langle r_i - c, \hat{r}_j \rangle \hat{r}_j}{\left\| (r_i - c) - \sum_{j=1}^{i-1} \langle r_i - c, \hat{r}_j \rangle \hat{r}_j \right\|} \quad (2)$$

where $\|x\|$ is shorthand for $\langle x, x \rangle^{1/2}$. In the following, we will repeatedly require the evaluation of $\langle \cdot, \hat{r}_i \rangle$ and $\|\pi(\cdot; \cdot)\|$. Although (1) and (2) can be evaluated explicitly every time, it can be more convenient to represent the (squared) norm after projection in terms of inner products (especially in kernel spaces):

$$\|\pi(x - c)\|^2 = \sum_{i=1}^k \langle x - c, \hat{r}_i \rangle^2 \quad (3)$$

since all \hat{r}_i are normalized and pairwise orthogonal. We can reduce $\langle \cdot, \hat{r}_i \rangle$ to

$$\langle x - c, \hat{r}_i \rangle = \frac{\langle c, c \rangle - \langle c, x \rangle - \langle c, r_i \rangle + \langle x, r_i \rangle - \sum_{j=1}^{i-1} \langle x - c, \hat{r}_j \rangle \langle r_i - c, \hat{r}_j \rangle}{(\langle c, c \rangle - 2\langle c, r_i \rangle + \langle r_i, r_i \rangle - \sum_{j=1}^{i-1} \langle r_i - c, \hat{r}_j \rangle^2)^{1/2}} \quad (4)$$

which can also be used recursively to compute the $\langle r_i - c, \hat{r}_j \rangle$ in (4). In the non-affine case, $c = \mathbf{0}$, (4) simplifies to

$$\langle x, \hat{r}_i \rangle = \frac{\langle x, r_i \rangle - \sum_{j=1}^{i-1} \langle x, \hat{r}_j \rangle \langle r_i, \hat{r}_j \rangle}{(\langle r_i, r_i \rangle - \sum_{j=1}^{i-1} \langle r_i, \hat{r}_j \rangle^2)^{1/2}} \quad (5)$$

Note that the denominator and parts of the nominator need to be computed just once. Further, we omit the explicit computation of any \hat{r}_i which would be infeasible in, e.g., RBF kernel and general inner product spaces. With dynamic programming, $\|\pi(x - c)\|^2$ can be computed in $\Theta(pk^2)$ time, where p is the effort required to compute an inner product.

In spatial indexing, pivots have been successfully used to bound distances via the triangle inequality [7, 8]. We propose to bound distances in terms of a decomposition of the squared Euclidean norm into dot products given by

$$d_{Euc}(x, y)^2 = \|x - y\|^2 = \langle x - y, x - y \rangle = \langle x, x \rangle + \langle y, y \rangle - 2\langle x, y \rangle \quad (6)$$

From this we can derive bounds for the Euclidean distance between two points given a bound on the dot product $\langle x, y \rangle$, assuming $\langle x, x \rangle$ and $\langle y, y \rangle$ are known. Let $\hat{r}_1, \dots, \hat{r}_k$ be pivot points previously orthogonalized by the Gram-Schmidt process as defined in Sect. 3. We can decompose $x - c$ and $y - c$ into k components aligned along the \hat{r}_i and one orthogonal remainder. We will call this $(k + 1)$ -th component x_\perp and y_\perp , respectively. It then follows that

$$\langle x - c, y - c \rangle = \langle x_\perp, y_\perp \rangle + \sum_{i=1}^k \langle \langle x - c, \hat{r}_i \rangle \hat{r}_i, \langle y - c, \hat{r}_i \rangle \hat{r}_i \rangle \quad (7)$$

Because the \hat{r}_i are pairwise orthogonal, this decomposition is uniquely defined. Since all \hat{r}_i have a unit norm, we can rewrite this equation to

$$\langle x, y \rangle = \langle x_\perp, y_\perp \rangle + \langle c, x \rangle + \langle c, y \rangle - \langle c, c \rangle + \sum_{i=1}^k \langle x - c, \hat{r}_i \rangle \langle y - c, \hat{r}_i \rangle \quad (8)$$

All of the terms on the right-hand side then either depend on x or y , but not on both, except for $\langle x_\perp, y_\perp \rangle$. In the semantics of Euclidean spaces, both x_\perp and y_\perp lie in the same $(d - k)$ -dimensional linear subspace. We can compute both as $x_\perp = (x - c) - \pi(x - c)$ and $y_\perp = (y - c) - \pi(y - c)$, respectively, but do not

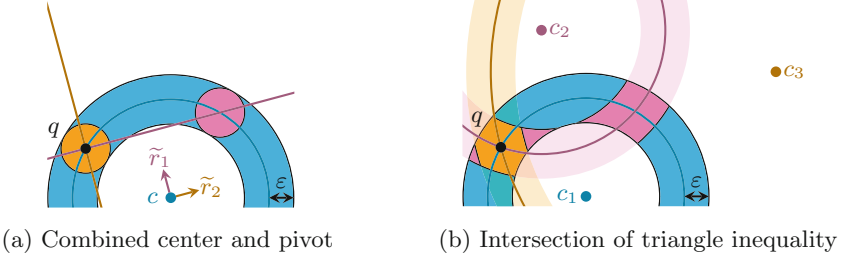


Fig. 1. Eligible search spaces around a query point q after filtering with the lower bounds obtained from **one**, **two**, or **three** centers and/or pivots.

know their relative orientation. Yet, we can bound their inner product using the Cauchy-Schwarz inequality resulting in the bounds $\pm(\langle x_{\perp}, x_{\perp} \rangle \cdot \langle y_{\perp}, y_{\perp} \rangle)^{1/2}$. By orthogonality of x_{\perp} and $\pi(x - c)$ we know $\|x_{\perp}\|^2 = \|x - c\|^2 - \|\pi(x - c)\|^2$. The bounds for the inner product $\langle x - c, y - c \rangle$ then follow as

$$\begin{aligned} & \langle c, x \rangle + \langle c, y \rangle - \langle c, c \rangle + \sum_{i=1}^k \langle x - c, \hat{r}_i \rangle \langle y - c, \hat{r}_i \rangle \\ & \pm \left(\left(\langle x, x \rangle + \langle c, c \rangle - 2 \langle c, x \rangle - \sum_{i=1}^k \langle x - c, \hat{r}_i \rangle^2 \right) \cdot \left(\langle y, y \rangle + \langle c, c \rangle - 2 \langle c, y \rangle - \sum_{i=1}^k \langle y - c, \hat{r}_i \rangle^2 \right) \right)^{1/2} \end{aligned} \quad (9)$$

which in the non-affine case, $c = \mathbf{0}$, becomes

$$\sum_{i=1}^k \langle x, \hat{r}_i \rangle \langle y, \hat{r}_i \rangle \pm \left(\left(\langle x, x \rangle - \sum_{i=1}^k \langle x, \hat{r}_i \rangle^2 \right) \cdot \left(\langle y, y \rangle - \sum_{i=1}^k \langle y, \hat{r}_i \rangle^2 \right) \right)^{1/2} \quad (10)$$

Inserting both of these values into (6) gives bounds on the squared Euclidean distance and, consequentially, on the Euclidean distance. These bounds are a generalization of at least two bounds known from the literature. When we assume the affine case and $k = 0$ pivots, the bounds derived from (6) and (10) reduce to

$$\langle x, x \rangle + \langle y, y \rangle - 2 \langle c, x \rangle - 2 \langle c, y \rangle + 2 \langle c, c \rangle \pm 2 \|x - c\| \|y - c\| \quad (11)$$

$$= (\|x - c\| \pm \|y - c\|)^2 \quad (12)$$

which are the bounds easily derivable from the triangle inequality. For the non-affine case with $k = 1$ pivots and normalized x and y , the inner product bounds (10) reduce to

$$\langle x, \hat{r}_1 \rangle \langle y, \hat{r}_1 \rangle \pm \left(\left(1 - \langle x, \hat{r}_1 \rangle^2 \right) \left(1 - \langle y, \hat{r}_1 \rangle^2 \right) \right)^{1/2} \quad (13)$$

which is the triangle inequality for cosines introduced in [10]. Triangle-inequality-based bounds have been used in spatial indexing in methods like, e.g., LAESA [7]. For multiple pivots, these approaches take the minimum or maximum of the

bounds obtained separately for each pivot. In our terminology, we refer to such pivots as centers c . Those are fundamentally different from the term pivots introduced here: When performing an ε -range query for a query point y , the eligible search space for vectors x according to the upper bound in (12) is a hyperspherical shell centered at c . This geometric shape can be described as the sumset (the set of all sums of pairs in the cartesian product) of a $(d-1)$ -sphere of radius $\|y-c\|$ centered at c and a d -ball of radius ε . When using pivots as per our definition, each pivot induces a hyperplane orthogonal to the \hat{r}_i which intersects with the hypersphere. Consequentially, the resulting eligible search space is the sumset of a $(d-1-k)$ -sphere of radius $(\|y-c\|^2 - \|\pi(y-c)\|^2)^{1/2}$ and a d -ball of radius ε . This is illustrated in two dimensions in Fig. 1. Each of the pivots eliminates an entire dimension from the sphere-part of the search space whereas the minimum lower bounds obtained from multiple centers produce an intersection of multiple hyperspherical shells. While $d-1$ pivots can reduce the search space to the sumset of at most 2 points and an ε -ball, the intersection of even d hyperspherical shells in the best case produces a volume that can be roughly described as a distorted hypercube with an “edge length” of about 2ε . The resulting volume can be exponentially larger in d than the search volume using $d-1$ pivots. As the volumes of regular shapes in Euclidean space expand exponentially in dimensions, one would expect an approximately exponential reduction in search space over an increasing number of pivots, whereas using the minimum upper bound over multiple centers does not induce such a reduction in search space volume. It is, therefore, of little surprise that the cosine bounds introduced in [10] ($k=1$), produced tighter bounds empirically than the triangle inequality ($k=0$), and were successfully applied to improve the performance of spherical k-means clustering [11]. Qualitatively, there is a clear argument for using a larger amount of pivots. However, the reduction in search space comes at the price of increased computational cost as the evaluation of $\langle y, \hat{r}_i \rangle$ is quadratic and the evaluation of the bounds is linear in k . Blindly increasing k is not universally advantageous for the computational cost of spatial indexing queries. But how many pivots tighten the bounds enough to counterweigh the overhead? More precisely, how much more of a point’s squared norm does the k -th randomly drawn pivot drawn cover on average? Although the answer does not refer to an optimal pivot choice, by arguing over expectations of underlying distributions, this conservative argument likely holds for previously unknown query points.

3 Expected Variance of Random Projections

The analysis of squared norms after projection is closely related to spectral analysis. If we chose any normalized vector v , $\mathbb{E}_{x \in X} [\|\pi(x - \mathbb{E}_{y \in X} [y]; v)\|^2]$ is simply the variance of X in direction v . Consequentially, for any pair of a normalized eigenvector e_i and the corresponding eigenvalue $\lambda_i^{(C(X))}$, we know that $\mathbb{E}_{x \in X} [\|\pi(x; e_i)\|^2] = \lambda_i^{(C(X))}$ for any origin-centered X . By orthogonality of the eigenvectors, this argument can be extended to any number of eigenvectors e_1, \dots, e_n as

$$\mathbb{E}_{x \in X} \left[\|\pi(x; e_1, \dots, e_n)\|^2 \right] = \sum_{i=1}^n \lambda_i^{(C(X))} \quad (14)$$

Pearson [9] showed that the eigenvectors of the covariance matrix are precisely the maximizers of this term, i.e. they are the solution to

$$\arg \max_{e_1, \dots, e_n} \mathbb{E}_{x \in X} \left[\|\pi(x; e_1, \dots, e_n)\|^2 \right] \quad (15)$$

If one intended to evaluate how much of the squared norm of any point is remaining after the projection onto k directions maximally, the answer immediately follows from the sum of the k largest eigenvalues. Employing the corresponding eigenvectors as \hat{r}_i would then be a reasonable approach. Yet, both eigenvectors and eigenvalues can be sensitive to noise in limited data sets [4]. They may not be an optimal choice when new and unknown data arises. We, hence, focus on the expectation of these values for a random set of reference points drawn from the data. More precisely we inspect

$$E_k^\Sigma(X) := \mathbb{E}_{\substack{r_1, \dots, r_k \in X \\ \forall i \neq j: r_i \neq r_j}} \left[\mathbb{E}_{x \in X} \left[\|\pi(x - c; r_1 - c, \dots, r_k - c)\|^2 \right] \right] \quad (16)$$

As with the eigenvectors and eigenvalues of the covariance matrix, this expected value is the sum of components introduced by each additional reference point taken into consideration. This naturally sums up the total variance of the data set for $k = d$. Through varying k we can obtain a cumulative description of how much variance an arbitrary linear projection within the data set can explain and the difference of neighboring values gives the amount of variance explained at random by the k -th component. We will write this difference as $E_k(X) := E_k^\Sigma(X) - E_{k-1}^\Sigma(X)$ where $E_0^\Sigma(X) = 0$. It follows that $E_k^\Sigma(X) = \sum_{i=1}^k E_i(X)$. Practically evaluating the expected value from any data set X for any $k \gg 1$ is infeasible, as it involves $\binom{|X|}{k}$ possible sets of reference points. It is much easier to estimate the value by the Monte Carlo method (i.e. choosing a fixed number of random sets of reference points) or to approximate it from the covariance matrix if it well describes the data set's distribution.

We will only consider the non-affine case of $c = \mathbf{0}$, as the affine case is analogous and introduces numerous subtractions hindering readability. We will also omit the constraint that the reference points must not be linearly dependent to improve readability. Starting from (16) we can deduce

$$E_k(X) = E_k^\Sigma(X) - E_{k-1}^\Sigma(X) = \mathbb{E}_{\substack{x \in X, \\ r_1, \dots, r_k \in \tilde{X}}} \left[\left\langle x, \frac{r_k - \pi(r_k; r_1, \dots, r_{k-1})}{\|r_k - \pi(r_k; r_1, \dots, r_{k-1})\|} \right\rangle^2 \right] \quad (17)$$

Here the term $r_k - \pi(r_k; r_1, \dots, r_{k-1})$ is the projection of r_k onto the linear subspace orthogonal to all r_1, \dots, r_{k-1} . We can represent this projection by a matrix multiplication with a matrix, which we will call A_{k-1} .

$$= \mathbb{E}_{\substack{x \in X, \\ r_1, \dots, r_k \in \tilde{X}}} \left[\frac{\langle x, A_{k-1} r_k \rangle^2}{\langle A_{k-1} r_k, A_{k-1} r_k \rangle} \right] = \mathbb{E}_{\substack{x \in X, \\ r_1, \dots, r_k \in \tilde{X}}} \left[x^T \frac{A_{k-1} r_k r_k^T A_{k-1}^T}{\text{tr}(A_{k-1} r_k r_k^T A_{k-1}^T)} x \right] \quad (18)$$

By rewriting $r_i r_i^T$ as R_i this further simplifies to

$$= \mathbb{E}_{\substack{x \in X, \\ r_1, \dots, r_k \in \tilde{X}}} \left[x^T \frac{A_{k-1} R_k A_{k-1}^T}{\text{tr}(A_{k-1} R_k A_{k-1}^T)} x \right] \quad (19)$$

$$= \text{tr} \left(\mathbb{E}_{r_1, \dots, r_{k-1} \in \tilde{X}} \left[\mathbb{E}_{r_k \in X} \left[\frac{A_{k-1} R_k A_{k-1}^T}{\text{tr}(A_{k-1} R_k A_{k-1}^T)} \right] \right] \mathbb{E}_{x \in X} [x x^T] \right) \quad (20)$$

By replacing $\mathbb{E}_{x \in X} [x x^T]$ with the covariance matrix $C(X)$ and renaming the innermost expected value to $C_k(X)$ we then obtain

$$= \mathbb{E}_{r_1, \dots, r_{k-1} \in \tilde{X}} [\text{tr}(C_k(X) C(X))] \quad (21)$$

A_0 is the identity matrix I_d , as the linear subspace orthogonal to an empty set of vectors is the entire space. Consequentially, we can define A_k recursively as

$$A_k = A_{k-1} - \frac{A_{k-1} R_k A_{k-1}^T}{\text{tr}(A_{k-1} R_k A_{k-1}^T)} = A_{k-1} - \frac{A_{k-1} R_k A_{k-1}}{\text{tr}(A_{k-1} R_k A_{k-1})} \quad (22)$$

As all R_i are symmetric, all A_i are symmetric as well. The expected value over r_k of $\frac{A_{k-1} R_k A_{k-1}}{\text{tr}(A_{k-1} R_k A_{k-1})}$ now (approximately) equals the covariance matrix of X after being projected to the linear subspace orthogonal to r_1, \dots, r_{k-1} and normalized. It follows immediately that $C_1(X) = C(\tilde{X})$ and thereby $E_1(X) = \text{tr}(C(\tilde{X}) C(X))$. However, $E_k(X)$ for $k > 1$ is much less easily defined because the A_i are dependent on the effective values of all r_j , $j \leq i$, and not only on r_i . To circumvent the problem we assume that all A_i are aggregate matrices just like $C(X)$ and sufficiently independent of each other to evaluate the $C_k(X)$ recursively. To highlight this assumption we will denote the approximated A_i as a function of X as $A_i(X)$. We further assume that all $A_i(X)$, $C_i(X)$, and $C(X)$ admit the same eigenvectors, whereby

$$E_k(X) = \mathbb{E}_{r_1, \dots, r_{k-1} \in \tilde{X}} [\text{tr}(C_k(X) C(X))] = \sum_{i=1}^d \lambda_i^{(C_k(X))} \lambda_i^{(C(X))} \quad (23)$$

We will hereafter omit the (X) in superscripts of eigenvalues for readability. Although the resulting values are no longer exact due to these two assumptions, they allow us to approximate the expected value by deriving the value of $\lambda_i^{(C_k)}$. Assuming that X is multivariate normally distributed, we can extract this value from the definition of $C_k(X)$ using the corresponding eigenvector e_i :

$$\lambda_i^{(C_k)} = e_i^T C_k(X) e_i = \text{tr} \left(e_i e_i^T C_k(X) \right) \quad (24)$$

$$= \mathbb{E}_{r_k \in X} \left[\frac{r_k^T A_{k-1}(X) e_i e_i^T A_{k-1}(X) r_k}{r_k^T A_{k-1}(X)^2 r_k} \right] \quad (25)$$

$$= \mathbb{E}_{r_k \in \mathcal{N}_{0_d, 1_d}} \left[\frac{r_k^T C(X)^{1/2} A_{k-1}(X) e_i e_i^T A_{k-1}(X) C(X)^{1/2} r_k}{r_k^T C(X)^{1/2} A_{k-1}(X)^2 C(X)^{1/2} r_k} \right] \quad (26)$$

$$= \mathbb{E}_{r_k \in \mathcal{N}_{0_d, 1_d}} \left[\frac{r_k^T e_i e_i^T C(X) A_{k-1}(X)^2 r_k}{r_k^T C(X) A_{k-1}(X)^2 r_k} \right] \quad (27)$$

We now substitute $C(X)A_{k-1}(X)^2$ with $D_{k-1}(X)$ which entails $\lambda_j^{(C)} \left(\lambda_j^{(A_{k-1})} \right)^2$ is equal to $\lambda_j^{(D_{k-1})}$. In favor of brevity we will omit the exponent (D_{k-1}) from here on. As per Proposition 2 in Kan and Bao [2], $\lambda_i^{(C_k)}$ then equals

$$= \int_0^\infty \frac{\text{tr} \left(e_i e_i^T D_{k-1}(X) (\mathbf{I}_d + 2tD_{k-1}(X))^{-1} \right)}{|\mathbf{I}_d + 2tD_{k-1}(X)|^{1/2}} dt \quad (28)$$

$$= \int_0^\infty \frac{\lambda_i}{(1 + 2t\lambda_i)^{1/2} \prod_{j=1}^d (1 + 2t\lambda_j)^{1/2}} dt \quad (29)$$

This integral is closely related to elliptic integrals and we do not provide a simple and closed-form solution. Solving the integral numerically would again involve too much computational effort. We instead propose to substitute the λ_j in the denominator with $(\lambda_i^2 \prod_{j=1}^d \lambda_j)^{1/(d+2)}$ whereby the integral takes the form of a scaled beta prime distribution:

$$\lambda_i^{(C_k)} \approx \lambda_i B(\alpha, \beta) \int_0^\infty \frac{t^{\alpha-1} (1 + 2(\lambda_i^2 \prod_{j=1}^d \lambda_j)^{\frac{1}{d+2}} t)^{-\alpha-\beta}}{B(\alpha, \beta)} dt \quad (30)$$

where $\alpha = 1$, $\beta = \frac{d}{2}$, and $B(\alpha, \beta)$ is the beta function. The integral over the scaled beta distribution is known to equal the scaling factor, whereby

$$\lambda_i^{(C_k)} \approx \frac{\lambda_i B(\alpha, \beta)}{2(\lambda_i^2 \prod_{j=1}^d \lambda_j)^{\frac{1}{d+2}}} \propto \lambda_i^{\frac{d}{d+2}} \quad (31)$$

As the $\lambda_i^{(C_k)}$ are eigenvalues of a normalized distribution, their sum must equal 1. Using this constraint, we can drop all factors independent of λ_i and derive

$$\lambda_i^{(C_k)} \approx \lambda_i^{\frac{d}{d+2}} / \sum_{j=1}^d \lambda_j^{\frac{d}{d+2}} \quad (32)$$

As the λ_j are dependent on $\lambda_j^{(C)}$ and $\lambda_j^{(A_{k-1})}$, this leads to the recursive definition

$$\lambda_i^{(C_k)} \approx \frac{(\lambda_i^{(C)} (\lambda_i^{(A_{k-1})})^2)^{\frac{d}{d+2}}}{\sum_{j=1}^d (\lambda_j^{(C)} (\lambda_j^{(A_{k-1})})^2)^{\frac{d}{d+2}}} \quad \lambda_i^{(A_k)} \approx \lambda_i^{(A_{k-1})} - \lambda_i^{(C_{k-1})} \quad (33)$$

This recursion terminates at $\lambda_i^{(A_0)} = 1$ and $\lambda_i^{(C_0)} = 0$. These approximations can be computed efficiently in $\Theta(dk)$ and inserted in (23) to give an approximation of $E_k(X)$. Since the approximations are based on the assumption that X is distributed according to some multivariate normal distribution they need not be accurate. Since all occurrences of any r_k in the formulae involve some sort of normalization, this approximation extends to any distribution of X for which $\{C(X)^{-1/2}x \mid x \in X\}$ is spherically symmetrically distributed, which includes cases like, e.g., d -balls. We also did not compensate for the requirement that all r_k must be pairwise different, as these arguments are based on distributions rather than point sets. In empirical tests the sample size, however, did not contribute to approximation quality. The biggest issue with this approximation is the fact, that while the A_i as variables in r_1 through r_i must have eigenvalues in

$\{0, 1\}$, the approximated eigenvalues $\lambda_i^{(A_k)}$ can become negative whereby latter E_k can be vastly overestimated. As we know that the $E_k^\Sigma(X)$ must sum to the total variance of X , we propose to cut off any excess in $E_k^\Sigma(X)$ and determine the $E_k(X)$ based on these cut values. To summarize, the approximation proceeds as follows: For all $1 \leq k \leq d$ compute the $\lambda_i^{(C_k)}$ values using the recursive formulations (33). Use these values to compute $E_i(X)$ values and reduce $E_i(X)$ values for larger k to not have their sum exceed the total variance of X , which compensates for negative $\lambda_i^{(A_k)}$. Even though this approximation from a theoretical point makes the wrong assumptions that the r_k are pairwise different and that the $C_i(X)$ are statistically independent, the approximation in our experiments gave close enough results to have it worth considering, especially as the exact computation of values has an enormous computational cost. The approximation via the Monte Carlo method is known to converge on the exact values, yet, might require enormous samples.

While (23) requires the covariance matrix of a mean-centered data set, the approach via Monte Carlo sampling applies directly to inner product values and, hence, to kernel spaces. The approximation in (23) can then be used in black-box optimization to obtain an approximate spectral analysis of the kernel space. The obtained spectrum is neglecting the scale of the eigenvalues of the covariance matrix as the $E_i(X)$ are invariant under the scaling of these values. In this manner, we can perform approximate spectral analysis even in spaces that do not allow for a direct approach, such as the RBF kernel space which has infinitely many dimensions. Naturally, the method must be applied in a truncated fashion for infinite dimensions, for which we here propose two solutions: Firstly, one can estimate $E_1(X)$ through $E_k(X)$ for some fixed k using the Monte Carlo method and rescale these values to sum to 1. This implies neglecting the remaining $d-k$ dimensions and assuming the data to have 0 variance along with these directions. The $d-k$ smallest eigenvalues of the covariance of such a data set must then be 0, too. Finding any set of k eigenvalues that leads to these $E_1(X)$ through $E_k(X)$ values then solves the truncated case. Secondly, one can assume that the remaining variance not explained by $E_k^\Sigma(X)$ is distributed over the remaining $d-k$ values according to some user-defined distribution. Assuming a uniform distribution, for example, would explain the remaining variance as noise in the embedding space which might be a reasonable assumption.

A special case can further be made on the evaluation of $E_k(X)$ values on normalized data. When working on \tilde{X} instead of X , which can be achieved in kernel space by dividing the occurrences of x in the formulae by $\langle x, x \rangle^{1/2}$, we immediately obtain that $E_1(\tilde{X})$ equals the sum of squared eigenvalues of $C(\tilde{X})$. While this equality does not hold for the approximation via eigenvalues of $C(\tilde{X})$, it is approximately obtained from the Monte Carlo method or precisely for an exhaustive evaluation of $E_1(\tilde{X})$. Just as the constraint of the sum of eigenvalues of $C(\tilde{X})$ equalling 1, this additional constraint can be used in the black-box optimization for retrieving the original eigenvalues from $E_k(\tilde{X})$ values. Using (31), these eigenvalues can be approximately translated into the relative eigenvalues of the non-normalized data whenever the data can be assumed to obey the distributional constraints of the approximation.

4 Random Projections and ID Estimation

As stated in the previous section, $E_1(\tilde{X})$ equals the sum of squared eigenvalues of $C(\tilde{X})$. The reciprocal of this specific value has been introduced as an estimator for intrinsic dimensionality named ABID [13], that is

$$\text{ID}_{ABID}(X) = E_1(\tilde{X})^{-1} = E_1^\Sigma(\tilde{X})^{-1} \quad (34)$$

For one, this observation adds additional semantics to the meaning of ABID as the number of basis vectors of a random projection to fully explain the variance in a data set. Yet, it also implies the applicability of the E_k values in the realm of ID estimation. Although E_1 gives the part of total variance a random projection based on in-distribution basis vectors can explain, not all E_k values are necessarily equal. That is, the projection onto two random directions does not necessarily cover twice the variance covered by projecting onto one random direction. This linearity is exclusively true for spherically symmetrical distributions such as d -balls and for all other distributions we would certainly expect $E_2^\Sigma(X) < 2E_1^\Sigma(X)$. Ultimately, we are looking for the smallest k such that $E_k^\Sigma(X) \geq \text{tr}(C(X))$, that is, the number of random projections required to explain the entire variance of X . Unfortunately, we only have formulae for integer k but we can generalize the approach of ABID in the sense of extrapolating from a fixed E_k which results in a parameterized ID estimator which we name the Thresholded Random In-distribution Projections (TRIP) Estimator:

$$\text{ID}_{TRIP}(X, k, \eta) = k + \frac{(1 - \eta) \text{tr}(C(X)) - E_k^\Sigma(X)}{E_k(X)} \quad (35)$$

where k is the number of considered projections and $\eta \in [0, 1]$ is a fraction describing how much of the variance we attribute to noise. Semantically this answers the question “*How many random projections are required to explain $(1-\eta)$ of the total variance if every further projection covers as much variance as the last one?*”. In the linear case of spherically symmetrical distributions as above, this estimator is ideally constant for $\eta = 0$ and all $1 \leq k \leq d$. On other distributions with $\eta = 0$, we would expect a curve that starts at (approximately, dependent on implementation) $\text{ID}_{ABID}(X)$ for $k = 1$ and approaches k for increasing k as the $E_i(X)$ are monotonically falling. Equality is likely only reached for $k = d$, as this requires zero variance after k projections, which is unlikely in presence of high-dimensional noise. The factor η is intended to compensate for this. For $\eta > 0$, the curve again starts at approximately $\text{ID}_{ABID}(X)$, approaches k , and after some k drops below it. As for parameter choice, η is application dependent whereas k can either be chosen empirically, or we can inspect values $1 \leq k \leq d$ to find the k at which $\text{ID}_{TRIP}(X, k, \eta)$ is closest to k . The latter is likely not feasible in a local ID fashion when using the Monte Carlo or exhaustive methods but can be done when using the approximation introduced in Sect. 3. When using a fixed k , obtaining an ID below this k is a strong indicator of having chosen k too large. In addition, the curve of $\text{ID}_{TRIP}(X, k, \eta)$ over varying k , just like the curve of $E_i(X)$, gives insights into the local distribution characteristics of the data set that goes beyond ID estimation. These

curves can theoretically help distinguish different subspaces, even when they share similar local ID.

Referring back to the discussions of indexing with linear projections in Sect. 2, we can now state a clear connection between indexing with random in-distribution pivots and intrinsic dimensionality measures. The $E_k^{\mathcal{D}}(X)$ values answer how much variance on average is covered by a set of k random pivots. The expected covered variance is – in an idealized case of, e.g., uniformly distributed hyperballs – reciprocally related to intrinsic dimensionality. This is most explicitly stated in the relation to ABID and gives rise to the TRIP estimator above. Using this geometric concept of ID estimation, we can argue on an on-average appropriate number of pivots in spatial indexing. In Sect. 2 we observed that the eligible search space for range queries when using k pivots is the sunset of a $(d - 1 - k)$ -sphere and an ε -ball. The radius of the hypersphere is equal to the norm of the component orthogonal to all pivots, and roughly describes how close the bounds derived in Sect. 2 are to the true distances. But there is a clear limit as to how much precision one needs in a finite data set. If this radius drops below the distance between nearest points, removing this slack from the distance estimates does not improve the discriminability. By choosing $\eta = \delta^2 / \text{tr}(C(X))$ where δ is the, e.g., mean/median/ p -percentile of nearest neighbor distances, we can use the TRIP estimator to evaluate just how many random projections exhaust the discriminative potential of pivoted indexing on average.

5 Pivot Filtering Linear Scan

For quality evaluation of the bounds as well as to validate the theoretical claims, we embed the bounds in a simple and easy-to-implement index. During the initialization, we choose k random pivots. As mentioned in Sect. 2, we pre-compute all parts of the equations that are independent of query points such as $\langle x, \hat{r}_i \rangle$ or the denominators in (4). Range and n -nearest neighbor queries were then implemented according to Algorithms 1 and 2. The algorithms are quite similar to LAESA [7] but do not require aggregation of multiple bounds as discussed in Sect. 2. Both algorithms are at least linear in $|X|$, which should be accounted for when comparing the performance with tree-based indices. Integrating the bounds into a tree-based index is a nearby extension but out of the scope of this paper. Both Algorithms 1 and 2 are trivially adaptable to search for the largest instead of the smallest distances. This index is also trivially adaptable to work on inner products instead of distances by exchanging the bounds. For our experiments, we implemented the index in the Rust language and called the functions from a Python wrapper to compare them to the cKDTree and BallTree implementations of SciPy [15]. The source code is publicly available at <https://github.com/eth42/pfls>. Using this very simple index we investigated the theoretical claims and the quality of the bounds. Figure 2 displays the results of applying the index to the MNIST training data set. All queries were 100-nearest-neighbor queries for 1000 query points drawn from the same data set. We performed 100 queries for each set of parameters and instantiated a new index for each query. As seen

Algorithm 1 n -nearest neighbor query for distances

```

function QUERY( $y \in \mathbb{R}^d, n \geq 1$ )
   $ls \leftarrow$  lower bounds of  $d(x, y)$  for all  $x \in X$  as per (6) and (10)
   $h \leftarrow$  empty max heap
  sort  $X$  by ascending  $ls[x]$ 
  for  $x \in X$  do
    if  $|h| < n$  or ( $ls[x] < h.max.key$  and  $d(x, y) < h.max.key$ ) then
      push  $x$  onto  $h$  with key  $d(x, y)$ 
      if  $|h| > n$  then remove entry with largest key from  $h$ 
      else if  $ls[x] \geq h.max.key$  then break
  return  $h$  as array/list

```

Algorithm 2 range query for distances

```

function QUERY-RANGE( $y \in \mathbb{R}^d, \varepsilon \in \mathbb{R}$ )
   $ls, hs \leftarrow$  lower and upper bounds of  $d(x, y)$  for all  $x \in X$  as per (6) and (10)
   $v \leftarrow$  empty list
  for  $x \in X$  do
    if  $ls[x] < \varepsilon$  and ( $hs[x] < \varepsilon$  or  $d(x, y) < \varepsilon$ ) then Push  $x$  into  $v$ 
  return  $v$ 

```

in Fig. 2a, the number of distance computations initially drops exponentially as we increase the number of pivots, which supports the theoretical claim that each pivot effectively eliminates one dimension from the data set and reduces the remaining search space exponentially. For increasing k , the descent in distance computations diminishes as the bounds become tight enough to sufficiently discriminate on neighboring points, and the query time eventually increases due to the cost of computing the bounds. In Sect. 4, we argued that the bounds only need to be as tight as to differentiate between nearest neighbors. To validate this claim, we investigated the ID_{TRIP} values using an η equal to the 10-percentile of squared 1-nearest-neighbor distances divided by the total variance of the distribution. The smallest k for which $ID_{TRIP}(X, k, \eta) \leq k$ is around 150 as can be seen in Fig. 2c. The minimum computation time in Fig. 2b is around 100 but the query time at $k = 150$ is not that much larger than at $k = 100$. The exact percentile is an educated guess and could be supported by inspecting the histogram of nearest-neighbor distances. Yet, the region of k that provides low query times is wide enough that rough estimates and educated guesses are likely to give good results. We conclude that ID_{TRIP} can be used to estimate a proper value for k by deriving η from a percentile of 1-nearest neighbor distances. To estimate a proper k efficiently, the approximation introduced in Sect. 3 can be used, which practically is sufficiently similar to the values obtained from Monte Carlo sampling as displayed in Fig. 2c. Lastly, we compared query times on HSV color histograms of the ALOI data set with varying numbers of dimensions [12]. The considered variants consist of 110250 instances with 27, 126, and 350 dimensions, respectively. As can be seen in Fig. 3 the query performance of our index is

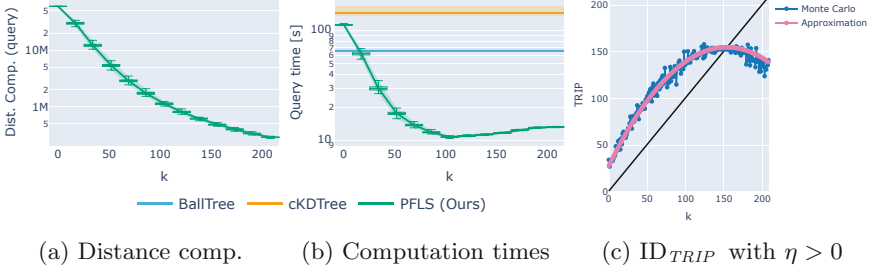


Fig. 2. Experimental results on varying numbers of pivots. Additional pivots exponentially reduce the distance computations, but the query time stagnates once the average discriminative power of the bounds has been exploited. A suitable number of pivots is suggested at the crossing point of ID_{TRIP} with the diagonal. Lines are average values, shaded area indicates the minimum and maximum.

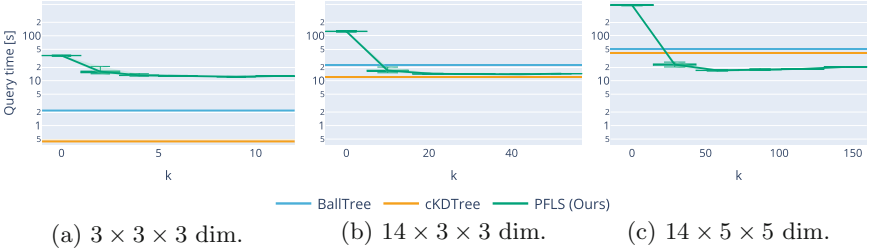


Fig. 3. Query times for ALOI color histograms with varying dimensionality.

mostly unaffected by increasing dimensionality. Due to our index using a linear scan, the tree-based reference implementations were faster on low dimensionality. For sufficiently high dimensional or small enough data sets, our index can outperform these reference implementations. For larger data sets, extending the approach to a tree-based structure appears promising.

6 Conclusion

In this paper, we introduced new bounds for Euclidean distances and inner products using a pivot-based approach. We showed that these bounds generalize the well-known bounds based on the triangle inequality. We argued why an increased number of pivots exponentially reduces the eligible search space of certain queries and derived an approach to estimate a reasonable number of pivots for practical purposes. We further showed how this number of pivots is intimately related to intrinsic dimensionality estimation. Lastly, we implemented the bounds in a simple and easily reproducible index that operates on both inner products and their induced distances and allows queries for the smallest and largest values. The empirical data presented aligns with the theoretical considerations and highlights the qualitative performance of implementing the bounds. Further research

should be invested in integrating these bounds into more sophisticated indices or constructing a tree-based index using these bounds.

References

1. Achtert, E., Böhm, C., Kriegel, H., Kröger, P., Zimek, A.: Robust, complete, and efficient correlation clustering. In: SIAM International Conference on Data Mining (SDM), pp. 413–418 (2007). <https://doi.org/10.1137/1.9781611972771.37>
2. Bao, Y., Kan, R.: On the moments of ratios of quadratic forms in normal random variables. *J. Multivar. Anal.* **117**, 229–245 (2013). <https://doi.org/10.1016/j.jmva.2013.03.002>
3. Chávez, E., Figueroa, K., Navarro, G.: Effective proximity retrieval by ordering permutations. *IEEE Trans. Pattern Anal. Mach. Intell.* **30**(9), 1647–1658 (2008). <https://doi.org/10.1109/TPAMI.2007.70815>
4. Everson, R.M., Roberts, S.J.: Inferring the eigenvalues of covariance matrices from limited, noisy data. *IEEE Trans. Signal Process.* **48**(7), 2083–2091 (2000). <https://doi.org/10.1109/78.847792>
5. Fukunaga, K., Olsen, D.R.: An algorithm for finding intrinsic dimensionality of data. *IEEE Trans. Comput.* **20**(2), 176–183 (1971). <https://doi.org/10.1109/T-C.1971.223208>
6. Houle, M.E., Kawarabayashi, K.: The effect of random projection on local intrinsic dimensionality. In: Reyes, N., et al. (eds.) SISAP 2021. LNCS, vol. 13058, pp. 201–214. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-89657-7_16
7. Micó, L., Oncina, J., Vidal, E.: A new version of the nearest-neighbour approximating and eliminating search algorithm (AESAs) with linear preprocessing time and memory requirements. *Pattern Recognit. Lett.* **15**(1), 9–17 (1994). [https://doi.org/10.1016/0167-8655\(94\)90095-7](https://doi.org/10.1016/0167-8655(94)90095-7)
8. Omohundro, S.M.: Five Balltree Construction Algorithms. International Computer Science Institute Berkeley, Berkeley (1989)
9. Pearson, K.: On lines and planes of closest fit to systems of points in space. *London, Edinb. Dublin Philos. Mag. J. Sci.* **2**(11), 559–572 (1901)
10. Schubert, E.: A triangle inequality for cosine similarity. In: Reyes, N., et al. (eds.) SISAP 2021. LNCS, vol. 13058, pp. 32–44. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-89657-7_3
11. Schubert, E., Lang, A., Feher, G.: Accelerating spherical k -means. In: Reyes, N., et al. (eds.) SISAP 2021. LNCS, vol. 13058, pp. 217–231. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-89657-7_17
12. Schubert, E., Zimek, A.: ELKI multi-view clustering data sets based on the Amsterdam library of object images (ALOI). Zenodo (2010). <https://doi.org/10.5281/zenodo.6355684>
13. Thordsen, E., Schubert, E.: ABID: angle based intrinsic dimensionality. In: Satoh, S., et al. (eds.) SISAP 2020. LNCS, vol. 12440, pp. 218–232. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-60936-8_17
14. Vadicamo, L., Gennaro, C., Amato, G.: On generalizing permutation-based representations for approximate search. In: Reyes, N., et al. (eds.) SISAP 2021. LNCS, vol. 13058, pp. 66–80. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-89657-7_6
15. Virtanen, P., et al.: SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat. Methods*, **17**, 261–272 (2020). <https://doi.org/10.1038/s41592-019-0686-2>