Principal Axis Transformation and SVD PhL 6.9.12

Motivation: this is a similar "active vs passive" topic to the vector discussion nearby, but the objects of interest here are matrices instead of vectors. I have never really gotten this subject written down. It occurs with the inertia tensor, the stress tensor, the Ng covariance tensor, and so on.

Standard tensor notation and graphical representation.

Use developmental tensor doc notation. Here then is how rank 1 and rank 2 tensors transform :

$V'_a = R_{a\alpha} V_\alpha$ // contravariant vector

$\mathbf{V}' = R\,\mathbf{V}$ // previous line in vector notation

$\mathbf{x}' = R\,\mathbf{x}$ // example of a vector transformation

$M'_{ab} = R_{a\alpha}\, R_{b\beta} M_{\alpha\beta}$ // contravariant rank 2 tensor

$M' = RMR^T = RMR^{-1}$ // previous line in matrix notation

So the above equations say that **V** and **x** are vectors with respect to transformation R, and that M is a rank-2 tensor with respect to R. Notice that the idea that **V**' = R **V** works with two of my three vector viewing pictures:

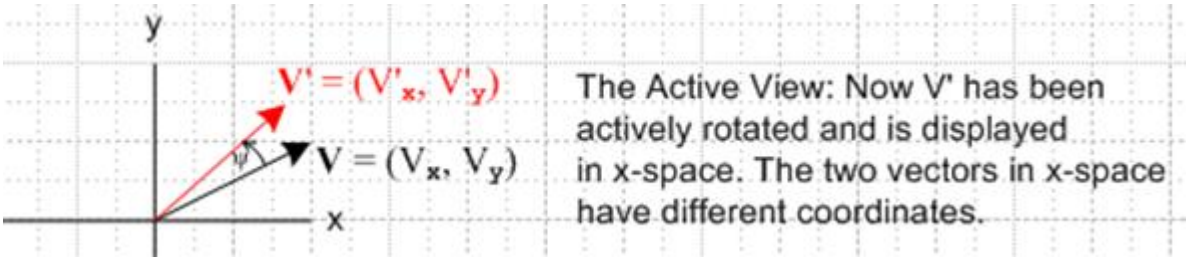

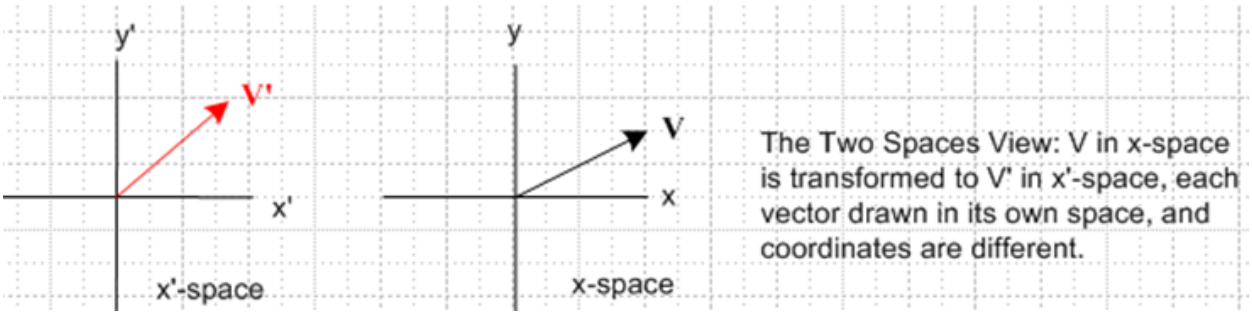

Now consider another set of equations

$\hat{\mathbf{u}}'^{(i)} = R^{-1}\,\hat{\mathbf{u}}^{(i)}$ // previous line applied to i^{th} basis vector $\mathbf{u}^{(i)}$

$\hat{u}'^{(i)}_a = R^{-1}_{ab}\, \hat{u}^{(i)}_b$ // previous line in components

Here we see that $\hat{\mathbf{u}}^{(i)}$ are vectors with respect to transformation R^{-1} (not R as for the vectors in the first list of equations above). If we multiply both sides of the last equation above by $\hat{u}^{(j)}_a$ we get

$$\hat{u}^{(j)}_a\, \hat{u}'^{(i)}_a = \hat{u}^{(j)}_a\, R^{-1}_{ab}\, \hat{u}^{(i)}_b = R^{-1}_{ji} = R_{ij}$$

But we also know that

$$\hat{u}^{(j)}_a\, \hat{u}'^{(i)}_a = \hat{\mathbf{u}}^{(j)} \bullet \hat{\mathbf{u}}'^{(i)}$$

and

$$\hat{u}^{(j)}_a\, \hat{u}'^{(i)}_a = \delta_{j,a} \hat{u}'^{(i)}_a = \hat{u}'^{(i)}_j$$

To summarize, we have

$$\hat{\mathbf{u}}^{(j)} \bullet \hat{\mathbf{u}}'^{(i)} = \hat{u}'^{(i)}_j = R_{ij}$$

The fact that $\hat{u}'^{(i)}_j = R_{ij}$ says that row i of matrix R_{ij} is in fact the vector $\hat{\mathbf{u}}'^{(i)}$, so one could write

$$R^{T} = [\hat{\mathbf{u}}'^{(1)}, \hat{\mathbf{u}}'^{(2)}, \ldots. \hat{\mathbf{u}}'^{(n)}] .$$

Once we have the above two sets of equations, we can explore things. One thing we find is that

$$\mathbf{V} = V_i\, \hat{\mathbf{u}}^{(i)} = V'_i\, \hat{\mathbf{u}}'^{(i)}$$

is consistent.

Proof: Write the right equality as

$$V_i\, \hat{u}^{(i)}_a = V'_i\, \hat{u}'^{(i)}_a$$

Then consider:

$$V'_a = R_{a\alpha} V_\alpha \qquad\qquad \hat{u}'^{(i)}_a = R^{-1}_{a\beta} \hat{u}^{(i)}_\beta$$

Then we find that

$$V'_i\, \hat{u}'^{(i)}_a = (R_{i\alpha} V_\alpha)\, (R^{-1}_{a\beta} \hat{u}^{(i)}_\beta) = (R_{i\alpha} V_\alpha)\, (R^{-1}_{a\beta} \delta_{i,\beta}) = (R_{i\alpha} V_\alpha)\, (R^{-1}_{ai})$$

$$= (R_{i\alpha} V_\alpha)\, (R^{T}_{ai}) = (R^{T}_{ai} R_{i\alpha}) V_\alpha = \delta_{a,\alpha}\, V_\alpha = V_a$$

$$V_i\, \hat{u}^{(i)}_a = V_i\, \delta_{i,a} = V_a$$

Now things are compatible with the passive view picture

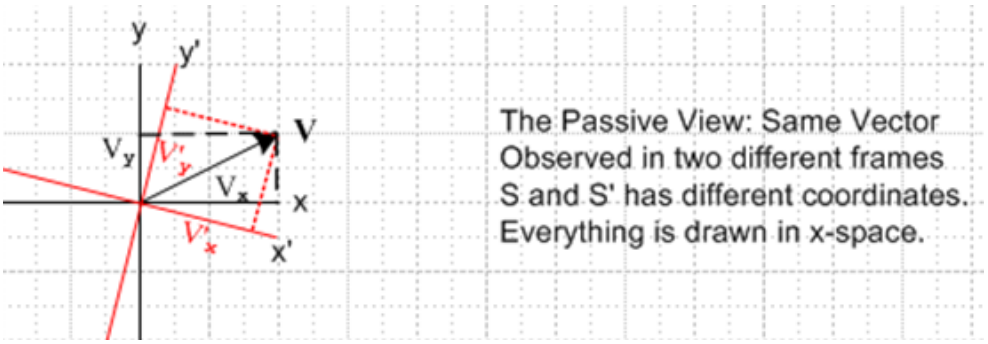

$\mathbf{V} = V_i \, \hat{\mathbf{u}}^{(i)} = V'_i \, \hat{\mathbf{u}}'^{(i)}$ // names of unit vectors don't match picture but OK

So when we "view" the vector **V** from the frame S', we find that it has certain components V'_i which are different from the components V_i that we "view" from frame S.

Focus on the rank 2 tensor.

Now how do we do this "view" stuff for a rank 2 tensor?

$M'_{ab} = R_{a\alpha} R_{b\beta} M_{\alpha\beta}$

Well, we can expand M in two different ways, following tensor doc

$M = M_{ij} \, \hat{\mathbf{u}}^{(i)} \, \hat{\mathbf{u}}^{(j)T}$ analogous to $\mathbf{V} = V_i \, \hat{\mathbf{u}}^{(i)}$

$M = M'_{ij} \, \hat{\mathbf{u}}'^{(i)} \, \hat{\mathbf{u}}'^{(j)T}$ analogous to $\mathbf{V} = V'_i \, \hat{\mathbf{u}}'^{(i)}$

I am sure that I could verify that both lines are consistent. Here we go

$M_{ab} = M_{ij} \, \hat{u}^{(i)}{}_a \, \hat{u}^{(j)T}{}_b = M_{ij} \, \delta_{i,a} \delta_{j,b} = M_{ab}$

$M'_{ij} \, \hat{u}'^{(i)}{}_a \, \hat{u}'^{(j)T}{}_b = (R_{i\alpha} R_{j\beta} M_{\alpha\beta}) \, (R^{-1}{}_{ac} \, \hat{u}^{(i)}{}_c) \, (R^{-1}{}_{bd} \, \hat{u}^{(j)}{}_d)$

$= (R_{i\alpha} R_{j\beta} M_{\alpha\beta}) \, (R^{-1}{}_{ac} \delta_{i,c}) \, (R^{-1}{}_{bd} \delta_{j,d})$

$= (R_{i\alpha} R_{j\beta} M_{\alpha\beta}) \, (R^{-1}{}_{ai}) \, (R^{-1}{}_{bj})$

$= (R^{-1}{}_{ai} R_{i\alpha}) \, (R^{-1}{}_{bj} R_{j\beta}) M_{\alpha\beta}) = M_{ab}$

Now how does the word "view" fit in here? The primed frame definitely has axis $\hat{\mathbf{u}}'^{(j)}$. And in the primed frame, we "see" components M'_{ij} which we can imagine to be diagonal with the right rotation.

Principal Axis Transformation (ie, diagonalization of M) and a theorem.

Imagine a rank-2 tensor M with components M_{ij} in reference frame S having unit vectors $\hat{\mathbf{u}}^{(i)}$. Imagine that R is an orthogonal transformation which brings M to diagonal form such that M' is diagonal.

$M' = RMR^{T}$ // previous line in matrix notation

$M'_{ab} = R_{a\alpha} R_{b\beta} M_{\alpha\beta}$ // contravariant rank 2 tensor

In reference frame S' the tensor M has components M'_{ij} and is diagonal, and this frame has basis vectors which are $\hat{\mathbf{u}}'^{(i)} = R^{-1} \hat{\mathbf{u}}^{(i)}$.

In reference frame S the tensor M has components M_{ij} and is not diagonal, and this frame has basis vectors which are $\hat{\mathbf{u}}^{(i)}$.

What can be said about certain M_{ij} elements? // $R^{-1}{}_{ij} = \hat{\mathbf{u}}'^{(i)}{}_{j}$

$M_{ab} = R^{-1}{}_{ac} R^{-1}{}_{bd} M'_{cd} = R^{-1}{}_{ac} R^{-1}{}_{bd} (\delta_{c,d} \lambda_c) = \lambda_c R^{-1}{}_{ac} R^{-1}{}_{bc}$

$= \lambda_c R^{-1}{}_{ac} R^{-1}{}_{bc} = \lambda_c \hat{\mathbf{u}}'^{(a)}{}_{c} \hat{\mathbf{u}}'^{(b)}{}_{c}$

That is to say

$M_{ab} = \Sigma_c \lambda_c \hat{u}'^{(a)}{}_{c} \hat{u}'^{(b)}{}_{c}$

$M_{aa} = \Sigma_c \lambda_c (\hat{u}'^{(a)}{}_{c})^2$

Imagine that the rotation R was selected such that the eigenvalues are in descending order

$\lambda_1 \geq \lambda_2 \ldots\ldots \geq \lambda_n$

and we make no claim yet about signs. If $\lambda_1 > 0$, then we may conclude that

$M_{aa} = \lambda_1 (\hat{u}'^{(a)}{}_{1})^2 \leq \lambda_1$ for all a = 1,2...n

I once thought that in the case all λ_i are positive, you could show that all the largest M_{aa} would be less than λ_1 and the second largest would be smaller than λ_2 and so on, but this is wrong as this counterexample shows:

```
> M := matrix(3,3,[1,1,1,1,2,2,1,2,3]);
```

$$M := \begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 2 \\ 1 & 2 & 3 \end{bmatrix}$$

```
> e := evalf(eigenvalues(M));
```

$$e := 5.048917340 + .2\ 10^{-10} I, .3079785280, .6431041320$$

```
> printf("%.2f    %.2f    %.2f",Re(e[1]),e[2],e[3]),
5.05    .31    .64
```

Here we have a real, square symmetric matrix M_{ij}, and we find the three eigenvalues. This matrix is seen to be positive definite. And we find that all diagonal $M_{ii} \le \lambda_1$ (= 5) as predicted. But we do not find that the second largest M_{ii} which here is 2 is less than λ_2 which is .64.

How unique is the R which does our principle axis transformation? I think the only freedom left is that you can use R or -R, and that must means reflecting all three of the primed unit vectors in S' space.

Theorem: Suppose the symmetric matrix M is positive definite. We know that

$$M_{aa} = \Sigma_c \lambda_c (\hat{u}'^{(a)}{}_c)^2$$

and therefore all the diagonal elements M_{aa} are positive. All we can say about these positive elements is that all of them are less than the largest eigenvalue λ_1.

Comments: Suppose you use this transformation on a galaxy of points as Ng does where M_{ij} is the covariance matrix. Then when you go to the new frame S', you can arrange for M'_{ij} to have descending positive eigenvalues and you know that M'_{11} will be larger than ALL the M_{ij} and that is all you really know. In frame S', you have found axes such that the variances step down monotonically, and you have no correlation between pairs of directions.

Ng Presentation

To make the Ng notation work, we have to make two notational changes.

(1) We need to make the replacement $R \rightarrow R^{-1}$ (same as $R \rightarrow R^T$) everywhere above. Among the results we find are

$\mathbf{V}' = R^{-1}\,\mathbf{V}$ // V is a rank 1 tensor (vector) with respect to R^{-1}

$M' = R^{-1}MR \quad = R^T MR$ // M is a rank 2 tensor with respect to R^{-1}

$\hat{\mathbf{u}}'^{(i)} = R\,\hat{\mathbf{u}}^{(i)}$ // each $\hat{\mathbf{u}}^{(i)}$ is a vector with respect to R

$\hat{\mathbf{u}}^{(j)} \bullet \hat{\mathbf{u}}'^{(i)} = \hat{u}'^{(i)}{}_j = R_{ji}$

Now after making this change we have the $\hat{\mathbf{u}}'^{(i)}$ being the columns of R

$R = [\hat{\mathbf{u}}'^{(1)}, \hat{\mathbf{u}}'^{(2)}, ... \hat{\mathbf{u}}'^{(n)}]$.

The form $M' = R^{-1}MR$ is the form one most often sees for a diagonalizing transformation, where the inverse is the matrix to the left of M.

(2) We change the names of the basis vectors as follows

$\hat{\mathbf{u}}^{(j)} \rightarrow \hat{\mathbf{x}}_j$ where $\hat{\mathbf{x}}_j$ is a unit vector along the jth feature axis

$\hat{\mathbf{u}}'^{(2)} \rightarrow \hat{\mathbf{u}}^{(j)}$ remove the prime.

Notice that the components of vector $\hat{\mathbf{x}}_j$ are $(\hat{\mathbf{x}}_j)_i$ and that $(\hat{\mathbf{x}}_j)_i = \delta_{j,i}$. Remember this just says that we have $\hat{\mathbf{x}}_1 = (1,0,0...0)$, $\hat{\mathbf{x}}_2 = (0,1,0...0)$, etc.

Then the above equation set becomes

$\mathbf{V}' = R^{-1}\,\mathbf{V}$ // V is a rank 1 tensor (vector) with respect to R^{-1}

$M' = R^{-1}MR = R^{T}MR$ // M is a rank 2 tensor with respect to R^{-1}

$\hat{\mathbf{u}}^{(i)} = R\,\hat{\mathbf{x}}_i$ // each $\hat{\mathbf{u}}^{(i)}$ is a vector with respect to R

$\hat{\mathbf{x}}_j \bullet \hat{\mathbf{u}}^{(i)} = \hat{u}^{(i)}{}_j = R_{ji}$

$R = [\hat{\mathbf{u}}^{(1)}, \hat{\mathbf{u}}^{(2)}, ... \hat{\mathbf{u}}^{(n)}]$

S = frame of reference having $\hat{\mathbf{x}}_i$ as the unit vectors

S_u = frame of reference having $\hat{\mathbf{u}}^{(i)}$ as the unit vectors

For Ng, M_{ij} is a positive definite covariance matrix. We rotate the basis vectors according to $\hat{\mathbf{u}}^{(i)} = R\,\hat{\mathbf{x}}_i$ and in that new frame of reference which we might call S_u we find that M has the diagonal form M'.

The SVD. Ng calls upon the "singular value decomposition" which in the most general case and with a standard notation says

$M = USV^{H}$ $\quad M = U\ S\ V^{H}$

$\quad\quad\quad

where we show the row x column sizes of the matrices. In general the elements of the matrices are complex numbers, and we have used ^{H} to indicate the Hermitian conjugate , as in

$V^{H} = (V^{T})^{*}$ or $V^{H}_{ij} = V_{ij}^{*}$

where * means complex conjugate. There are several notations for complex conjugate and for Hermitian conjugate and they conflict somewhat with one another

complex conjugate	use *; use overbar
Hermitian conjugate	use † ; use ^{H} ; use * ; use $^{+}$

In this decomposition, the square matrices U and V are both unitary, which means $U^{H} = U$ and $V^{H} = V$. The non-square matrix S is all zeros except for some real numbers on its "diagonal" which starts in the upper left corner. These numbers are real and are called "the singular values" of matrix M and they are conventionally listed in monotonic decreasing order.

If we write the U and V matrices in terms of their columns in this manner

$U = [\hat{u}^{(1)}, \hat{u}^{(2)}, ... \hat{u}^{(n)}]$

$V = [\hat{v}^{(1)}, \hat{v}^{(2)}, ... \hat{v}^{(n)}]$

then the $\hat{u}^{(i)}$ are eigenvectors of MM^{H} and the $\hat{v}^{(i)}$ are eigenvectors of $M^{H}M$. Notice that both these matrices MM^{H} and $M^{H}M$ are Hermitian ($A=A^{H}$) and therefore *have* eigenvectors. For real numbers only, the condition is that $A = A^{T}$ for a matrix to have eigenvectors. The claim is that the non-zero singular values (on the diagonal of S) are the square roots of the non-zero eigenvalues of MM^{H} and $M^{H}M$.

Our application of the SVD decomposition is much simpler than this general case. First of all, we have only real numbers so $A^{H} = A^{T}$ for any of our matrices. Secondly, dimension p = n, so all matrices are square n x n. In this case, then, the SVD says

$M = U S V^{T}$ where U and V are symmetric and S is diagonal.

In Ng's application of SVD, we identify U = V = R where R is real orthogonal so $R^{-1} = R^{T}$, and we identify S with M', our diagonal matrix, so the decomposition is then

$M = U S V^{T} = R\ M' R^{T}$ since $M' = R^{T}M R$ (see above)

Since M is symmetric, it has some eigenvectors $\hat{u}^{(i)}$ and eigenvalues λ_i. Then we can verify the claim made above that the diagonal elements of S should be the square roots of the eigenvalues of MM^{T} and $M^{T}M$. Of course these two matrices are the same and are equal to M^{2} so:

$MM^{T}\hat{u}^{(i)} = M^{T}M\,\hat{u}^{(i)} = M^{2}\hat{u}^{(i)} = M\lambda_i\hat{u}^{(i)} = \lambda_i M\hat{u}^{(i)} = \lambda_i\lambda_i\hat{u}^{(i)} = \lambda_i^{2}\hat{u}^{(i)}$.

If M is a positive definite matrix, then the descending elements of S = M' are the eigenvalues of M,

$$\lambda_1 \geq \lambda_2 ..... \geq \lambda_n > 0$$

and the eigenvectors which are the columns of U = R are the corresponding eigenvectors.

Main Point: We are trying to find the rotation matrix R which diagonalizes M into M'

$$M' = R^T M R$$

such that the diagonal elements of M' are in descending order. By calling the SVD routine with M as argument and asking for three returned matrices

$$[U, S, V] = \text{svd}(M) \qquad \text{meaning } M = U S V^T = R M' R^T$$

when U is our desired matrix R, and S is our desired matrix M'.

For the matrix M = Σ, the covariance matrix, we have

$$[U, S, V] = \text{svd}(\Sigma) \qquad \text{meaning } \Sigma = U \Sigma' V^T = R \Sigma' R^T$$

where Σ' is the covariance matrix observed in frame S_u which is diagonal.

Comment: Before diagonalization, a positive definite matrix M has all $M_{ii} > 0$. After diagonalization the new diagonal M' matrix has diagonal elements $M'_{ii} = \lambda_i$. These are all positive, and we know that the first one has the property that $M'_{11} \geq M_{ii}$ for all i. Secondly, we know that $M'_{ij} = 0$ when i and j are different, and this says that when viewed from frame S_u, there is no correlation between any pair of different directions.

Comment: Since the variances M'_{ii} decrease as i increases, we want to consider dropping dimensions at the end of our list, not the beginning of the list.

Reduction Idea: Ng says to do this for all the training samples $\mathbf{x}^{(i)}$

$$\underset{k}{\mathbf{z}^{(i)}} = \underset{k \text{ x } n}{(U_{reduced})^T} \underset{n}{\mathbf{x}^{(i)}}$$

[Note conflict now, $\mathbf{x}^{(i)}$ is a training sample, while $\hat{\mathbf{x}}^{(i)}$ is an axis in feature space.] Well here is what the means. We could start by doing this

$$\mathbf{x}'^{(i)} = U^T \mathbf{x}^{(i)} = R^{-1} \mathbf{x}^{(i)}$$

and $\mathbf{x}'^{(i)}$ is then our point $\mathbf{x}^{(i)}$ transformed to x'-space, so to speak, according to $\mathbf{V}' = R^{-1} \mathbf{V}$ in the Ng convention described above for R. The first k components of each $\mathbf{x}'^{(i)}$ form a k-dimensional vector

which Ng calls $\mathbf{z}^{(i)}$. Then we get exactly the above equation. This is the compressed feature data! Only the dimensions with the least variance were thrown out in this compression.

We know that

$$\mathbf{x}^{(i)} = U\, \mathbf{x}'^{(i)}$$

If we replace U by $U_{reduced}$, we get some approximation to the original data

$$\underset{n}{\mathbf{x}^{(i)}_{approx}} = \underset{n \times n}{U_{reduced}}\, \underset{n}{\mathbf{x}'^{(i)}} = \underset{n \times k}{U_{reduced}}\, \underset{k}{\mathbf{z}^{(i)}}$$

How much "variance" was lost in our process. Here is the measure

$$\text{fraction lost} = (1/m)\, \Sigma_i\, |\mathbf{x}^{(i)} - \mathbf{x}^{(i)}_{approx}|^2 \;/\; (1/m)\Sigma_i\, |\mathbf{x}^{(i)}|^2$$

$$= (1/m)\, \Sigma_i\, |\mathbf{x}^{(i)} - \mathbf{x}^{(i)}_{approx}|^2 \;/\; \Sigma_j\, \mathrm{var}(X_j)$$

where the denominator is the sum of the variances over all the dimensions of feature space, and the numerator is the sum of the variance error in all dimensions caused by the imperfect reconstruction.

Now how do we show that

$$(1/m)\, \Sigma_i\, |\mathbf{x}^{(i)} - \mathbf{x}^{(i)}_{approx}|^2 \;/\; (1/m)\Sigma_i\, |\mathbf{x}^{(i)}|^2 \;=\; 1 - (\,\Sigma_{i=1}^{k}\lambda_i / \Sigma_{i=1}^{n}\lambda_i)$$

I haven't done much with variances with vectors inside, but see you can just add components.

I do know that the trace of M is same as that of M', something I never used!

$$M_{11} + M_{22} + ... + M_{nn} = \lambda_1 + \lambda_2 + ... + \lambda_n$$

Well, here is a shot at it.

$$\mathbf{x}'^{(i)} = U^T\, \mathbf{x}^{(i)} = R^{-1}\, \mathbf{x}^{(i)}$$

$$\mathbf{x}^{(i)}_{approx} = R_r\, \mathbf{x}'^{(i)} = R_r\, R^{-1}\, \mathbf{x}^{(i)}$$

where

$$R = [\hat{\mathbf{u}}^{(1)}, \hat{\mathbf{u}}^{(2)}, ... \hat{\mathbf{u}}^{(n)}]$$

$$R_r = [\hat{\mathbf{u}}^{(1)}, \hat{\mathbf{u}}^{(2)}, ... \hat{\mathbf{u}}^{(k)}, 0,0,0..0]$$

I think that

$$R_r R^T = (\hat{\mathbf{u}}^{(1)})\, (\hat{\mathbf{u}}^{(1)})^T + (\hat{\mathbf{u}}^{(2)})\, (\hat{\mathbf{u}}^{(2)})^T + ... + (\hat{\mathbf{u}}^{(k)})\, (\hat{\mathbf{u}}^{(k)})^T$$

Then we have

$$\mathbf{x}^{(i)} - \mathbf{x}^{(i)}_{approx} = \mathbf{x}^{(i)} - R_r R^{-1} \mathbf{x}^{(i)} = (1 - R_r R^{-1}) \mathbf{x}^{(i)}$$

$$= (1 - \{(\hat{\mathbf{u}}^{(1)}) (\hat{\mathbf{u}}^{(1)})^T + (\hat{\mathbf{u}}^{(2)}) (\hat{\mathbf{u}}^{(2)})^T + ... + (\hat{\mathbf{u}}^{(k)}) (\hat{\mathbf{u}}^{(k)})^T \}) \mathbf{x}^{(i)}$$

$$= \Sigma_{s=k+1}^{n} (\hat{\mathbf{u}}^{(s)}) (\hat{\mathbf{u}}^{(s)})^T \mathbf{x}^{(i)}$$

Then

$$(1/m) \Sigma_i | \mathbf{x}^{(i)} - \mathbf{x}^{(i)}_{approx}|^2 = (1/m) \Sigma_i | (1 - R_r R^{-1}) \mathbf{x}^{(i)}|^2$$

$$= (1/m) \Sigma_i (1 - R_r R^{-1}) \mathbf{x}^{(i)} \bullet (1 - R_r R^{-1}) \mathbf{x}^{(i)}$$

Back up. Go back to

$M = M_{ij} \hat{\mathbf{u}}^{(i)} \hat{\mathbf{u}}^{(j)T}$ analogous to $\mathbf{V} = V_i \hat{\mathbf{u}}^{(i)}$

$M = M'_{ij} \hat{\mathbf{u}}'^{(i)} \hat{\mathbf{u}}'^{(j)T}$ analogous to $\mathbf{V} = V'_i \hat{\mathbf{u}}'^{(i)}$

and write these in Ng notation as

$$M = \Sigma_{i,j=1}^{n} M_{ij} \hat{\mathbf{e}}^{(i)} \hat{\mathbf{e}}^{(j)T}$$

$M = \Sigma_{i,j=1}^{n} M'_{ij} \hat{\mathbf{u}}^{(i)} \hat{\mathbf{u}}^{(j)T}$ where $\hat{\mathbf{u}}^{(i)} = R \hat{\mathbf{e}}^{(i)}$

Now I think he is setting the higher order u to zero, so we can say

$$M \approx \Sigma_{i,j=1}^{k} M'_{ij} \hat{\mathbf{u}}^{(i)} \hat{\mathbf{u}}^{(j)T} = \Sigma_{i=1}^{k} M'_{ii} \hat{\mathbf{u}}^{(i)} \hat{\mathbf{u}}^{(i)T}$$

So with full M, we have total variances being $\Sigma_{s=1}^{n} \lambda_s$ and for approximate M we have $\Sigma_{s=1}^{k} \lambda_s$, so this gives the notion that

variance retained = $\Sigma_{s=1}^{k} \lambda_s / \Sigma_{s=1}^{n} \lambda_s$

variance lost = $1 - \Sigma_{s=1}^{k} \lambda_s / \Sigma_{s=1}^{n} \lambda_s$

Now my only issue is how do you relate this to the other form. Well, I see now that

$\Sigma_{s=1}^{n} \lambda_s = \Sigma_{s=1}^{n} M_{ss}$ since trace is preserved

so you can regard either sum as the "total variance" as measured in either reference frame.

So the question remaining is this: why does

$$(1/m)\,\Sigma_i\,|\mathbf{x}^{(i)} - \mathbf{x}^{(i)}_{approx}|^2$$

represent the variance lost? It is reasonable, but I want a proof that it agrees with the λ stuff. Here is what I want to prove:

$$(1/m)\,\Sigma_i\,|\mathbf{x}^{(i)} - \mathbf{x}^{(i)}_{approx}|^2 \;/\; \Sigma_{s=1}^{n}\lambda_s \;=\; 1 - \Sigma_{s=1}^{k}\lambda_s \;/\; \Sigma_{s=1}^{n}\lambda_s$$

or

$$(1/m)\,\Sigma_i\,|\mathbf{x}^{(i)} - \mathbf{x}^{(i)}_{approx}|^2 \;=\; \Sigma_{s=1}^{n}\lambda_s \;-\; \Sigma_{s=1}^{k}\lambda_s \;=\; \Sigma_{s=k+1}^{n}\lambda_s$$

I did show above that

$$\mathbf{x}^{(i)} - \mathbf{x}^{(i)}_{approx} \;=\; \Sigma_{s=k+1}^{n}\,(\hat{\mathbf{u}}^{(s)})\,(\hat{\mathbf{u}}^{(s)})^{T}\,\mathbf{x}^{(i)}$$

but I don't see any nice way to use this. How about saying the variance is independent of rotated basis, so we have

$$(1/m)\,\Sigma_{i=1}^{n}\,|\mathbf{x}^{(i)} - \mathbf{x}^{(i)}_{approx}|^2 \;=\; (1/m)\,\Sigma_{i=1}^{n}\,|\mathbf{x}'^{(i)} - \mathbf{x}'^{(i)}_{approx}|^2$$

$$= (1/m)\,\Sigma_{i=1}^{n}\,\Sigma_{s=1}^{n}(\mathbf{x}'^{(i)}_s - \mathbf{x}'^{(i)}_{approx,s})^2$$

$$= (1/m)\,\Sigma_{i=1}^{n}\,\Sigma_{s=k+1}^{n}(\mathbf{x}'^{(i)}_s - \mathbf{x}'^{(i)}_{approx,s})^2$$

$$= (1/m)\,\Sigma_{i=1}^{n}\,\Sigma_{s=k+1}^{n}(\mathbf{x}'^{(i)}_s)^2$$

$$= \Sigma_{s=k+1}^{n}\,(1/m)\,\Sigma_{i=1}^{n}(\mathbf{x}'^{(i)}_s)^2$$

I need to show that

$$\Sigma_{s=k+1}^{n}\,(1/m)\,\Sigma_{i=1}^{n}(x'^{(i)}_s)^2 \;\;= \Sigma_{s=k+1}^{n}\,\lambda_s$$

which means I need to show that

$$(1/m)\,\Sigma_{i=1}^{n}(x'^{(i)}_s)^2 \;= \lambda_s$$

But this is precisely the variance in the sth component direction. QED