Polar Decomposition and the Closest Rotation

Julian Panetta

February 23, 2015

1 Polar Decomposition

A square $n \times n$ complex matrix A has two polar decompositions: the standard right polar decomposition $A = UP$ and the left polar decomposition $A = P_\ell U$ (also known as the reverse polar decomposition). Here U is a unitary matrix, and P and P_ℓ are positive semidefinite matrices. Intuitively, the right polar decomposition expresses any linear map as a stretching along n orthogonal axes followed by a norm-preserving transformation—e.g. a rotation/reflection. In the left polar decomposition, these steps' order is reversed.

1.1 Existence

The polar decomposition can be found using the SVD (though there are also numerical methods to compute it directly), which exists for all A:

$$A = W\Sigma V^*,$$

where W and V are unitary and Σ is a real, diagonal matrix with nonnegative entries. Inserting the $n \times n$ identity in the form of V^*V or W^*W we see:

$$A = WV^*V\Sigma V^* := UP, \quad A = W\Sigma W^*WV^* := P_\ell U,$$

with $U = WV^*$, $P = V\Sigma V^*$, and $P_\ell = W\Sigma W^*$.

1.2 Uniqueness (for Nonsingular A)

Assume another right polar decomposition exists (uniqueness of the left decomposition is proved analogously):

$$A = UP = U_2 P_2.$$

"Squaring" both sides:

$$A^*A = P^*P = P_2^*P_2$$

Taking the square root of both sides, which is unique because P and P_2 are positive semidefinite, this implies:

$$P = P_2$$

So even if A is singular, the stretching part is unique. If A is nonsingular, P must also be, and

$$UP = U_2 P \quad \implies U = U_2,$$

proving the full polar decomposition is unique.

However, if A is singular, then P must also be singular. In this case, an arbitrary rotation/reflection can be applied to P's nullvectors by U, provided U keeps them orthogonal to vectors from P's column space. So U is not unique in this case (though its action on P's column space still is).

2 Real Matrix Case

When A is a real matrix, matrices W and V of the SVD are also real, making U, P, and P_ℓ all real. This means the unitary transformation U is actually an orthogonal transformation.

Since $\det(P) \geq 0$,

$$\det(U) = \operatorname{sgn}(\det(A))$$

when A is nonsingular. This means, U is a rotation when A's determinant is positive and a reflection when it is negative.

3 Closest Unitary Matrix

We prove that the polar decomposition gives us the closest unitary matrix to A in the Frobenius norm sense:

$$\|A - U\|_F^2 < \|A - X\|_F^2 \quad \forall\, X \in U(n),\ X \neq U.$$

Furthermore, this shortest distance is actually

$$\|A - U\|_F^2 = \sum_i (\sigma_i - 1)^2,$$

where σ_i are the singular values of A.

3.1 Proof

First note:

$$\|A - U\|_F^2 = \operatorname{tr}((A-U)^*(A-U)) = \|A\|_F^2 - 2\operatorname{tr}(A^*U) + \operatorname{tr}(I) = \|A\|_F^2 - 2\operatorname{tr}(P) + \operatorname{tr}(I), \tag{1}$$

since $A^*U = P$ by the (right) polar decomposition. Likewise:

$$\|A - X\|_F^2 = \|A\|_F^2 - 2\operatorname{tr}(A^*X) + \operatorname{tr}(I) \tag{2}$$

Subtracting (2) from (1):

$$\|A - U\|_F^2 - \|A - X\|_F^2 = 2\operatorname{tr}(A^*X - P)$$

Using the eigenvalue decomposition $P = V\Lambda V^*$:

$$\begin{aligned}
\|A - U\|_F^2 - \|A - X\|_F^2 &= 2\operatorname{tr}(V\Lambda V^*U^*X - V\Lambda V^*) \\
&= 2\operatorname{tr}(\Lambda V^*U^*XV - \Lambda V^*V) = 2\operatorname{tr}(\Lambda Y - \Lambda) \\
&= 2\sum_i \lambda_i(Y_{ii} - 1) = 2\sum_i \lambda_i(\operatorname{Re}(Y_{ii}) - 1)
\end{aligned}$$

where $Y = V^*U^*XV$ is some unitary matrix, meaning $\operatorname{Re}(Y_{ii}) \leq 1$. Since P is positive semidefinite ($\lambda_i \geq 0$), this sum obviously obtains its maximum (of zero) only when $\operatorname{Re}(Y_{ii}) = 1$:

$$\|A - U\|_F^2 - \|A - X\|_F^2 \leq 0 \quad \Longrightarrow \quad \|A - U\|_F^2 \leq \|A - X\|_F^2.$$

Moreover, equality happens only when

$$\begin{aligned}
Y = V^*U^*XV &= I \\
U^*X = VV^* &= I \\
X &= U,
\end{aligned}$$

proving the inequality.

Finally,

$$\|A - U\|_F^2 = \|A\|_F^2 - 2\,\mathrm{tr}(P) + \mathrm{tr}(I) = \sum_i \sigma_i^2 - 2\lambda_i + 1$$

but notice

$$A^*A = P^*P \quad\Longrightarrow\quad V\Sigma^2V^* = V\Lambda^2V^* \quad\Longrightarrow\quad \lambda_i = \pm\sigma_i \tag{3}$$

by uniqueness of eigenvalues. So, since P is positive semidefinite, $\lambda_i = \sigma_i$, and we have:

$$\|A - U\|_F^2 = \sum_i \sigma_i^2 - 2\sigma_i + 1 = \sum_i (\sigma_i - 1)^2$$

4 Closest Rotation

In the case that A is real and $\det(A) > 0$, the polar decomposition U is a rotation, and it is therefore the closest rotation to A. When $\det(A) = 0$, an SVD can be computed such that WV^T is a rotation (the sign for the left and right nullspace basis vectors can be chosen arbitrarily). However, when $\det(A) < 0$, U must be a reflection. For U to be a rotation when $\det(A) < 0$, P would need to have an odd number of negative eigenvalues. This is easy enough to do—simply flip some signs in the SVD used to derive $A = UP$—but unfortunately our proof from 3.1 no longer applies. It's unclear how to maximize over Y_{ii} under the constraint that Y is a rotation. In particular, it's not obvious that the optimal Y_{ii} must all be ± 1 (though it turns out that they still are).

We resort to finding the closest rotation by a constrained minimization:

$$\min_{R\in\mathbb{R}^{n\times n}} \|A - R\|_F^2 \quad \text{s.t.} \quad R^TR = I,\ \det(R) = 1.$$

The corresponding Lagrangian is:

$$\mathcal{L}(R, M, m) = \|A - R\|_F^2 + \mathrm{tr}((R^TR - I)^TM) + m(\det(R) - 1),$$

where m is a single Lagrange multiplier, and M is a symmetric matrix of Lagrange multipliers (one per unique equation in the symmetric $R^TR = I$ constraint). Differentiating with respect to M and m just recovers the constraints, but differentiating with respect to R gives:

$$\frac{\partial\mathcal{L}}{\partial R} = -2\frac{\partial\,\mathrm{tr}(A^TR)}{\partial R} + \frac{\partial\,\mathrm{tr}((R^TR)M)}{\partial R} + m\frac{\partial\det(R)}{\partial R} = -2A + 2RM + m\det(R)[R^{-1}]^T \stackrel{!}{=} 0.$$

The differentiation rules used here are proved in Appendix A.

Since the constraints must be satisfied at the optimum, we can plug in $\det(R) = 1$ and $[R^{-1}]^T = R$:

$$2A = 2RM + mR \quad\Longrightarrow\quad A = R\left(M + \frac{m}{2}I\right).$$

This means that an optimal rotation R puts A in a very similar form to the polar decomposition: $A = RB$, where $B = \left(M + \frac{m}{2}I\right)$ is a symmetric matrix. Using eigendecomposition $B = Q\Lambda Q^T$,

$$A^TA = B^TB = Q\Lambda^2Q^T.$$

As in (3), $\lambda_i = \pm\sigma_i$, with the particular signs determined by R. Encoding the signs in diagonal matrix S so that $\lambda = S_{ii}\sigma_i$, we see that the optimal distance must be of the form:

$$\|A - R\|_F^2 = \|RQ\Lambda Q^T - R\|_F^2 = \mathrm{tr}(\Lambda^2 - 2\Lambda + I) = \sum_i(\lambda_i - 1)^2 = \sum_i(S_{ii}\sigma_i - 1)^2. \tag{4}$$

Furthermore, any assignment of signs multiplying to -1 is obtained by some rotation R. To see this, consider arbitrary signs, $\Lambda = S\Sigma$, $\det(S) = -1$:

$$A = W\Sigma V^T = W(SV^TVS)\Sigma V^T = \underbrace{(WSV^T)}_{R}\underbrace{(V\Lambda V^T)}_{B}, \tag{5}$$

and $\det(R) = \det(A)/\det(B) = 1$. In other words, we can find an optimal R by optimizing over $S_{ii} = \pm 1$.

Since $\sigma_i > 0$, each negative sign increases distance (4) by $4\sigma_i$, and the optimal sign assignment must negate only the single smallest singular value. Choosing these optimal signs, $S_{ii} = \begin{cases} 1 & i < n \\ -1 & i = n \end{cases}$ (assuming singular values are sorted in decreasing order), (5) gives the closest rotation $R = WSV^T$.

Appendix A Matrix Derivatives

Here we prove the three matrix derivatives we used on the Lagrangian.

First

$$\left[\frac{\partial \operatorname{tr}(A^T R)}{\partial R}\right]_{ij} = \frac{\partial (A_{kl} R_{kl})}{\partial R_{ij}} = A_{kl}\delta_{ik}\delta_{jl} = A_{ij}.$$

Second,

$$\left[\frac{\partial \operatorname{tr}((R^T R)M)}{\partial R}\right]_{ij} = \frac{\partial (R_{kl} R_{km} M_{ml})}{\partial R_{ij}} = (\delta_{ki}\delta_{lj} R_{km} + R_{kl}\delta_{ki}\delta_{mj}) M_{ml} = R_{im} M_{mj} + R_{il} M_{jl} = [R(M + M^T)]_{ij},$$

so when M is symmetric, $\frac{\partial \operatorname{tr}((R^T R)M)}{\partial R} = 2RM$.

Finally,

$$\left[\frac{\mathrm{d} \det(R)}{\mathrm{d} R}\right]_{ij} = \left.\frac{\mathrm{d}}{\mathrm{d}\epsilon}\right|_{\epsilon=0} \det(R + \epsilon e_i \otimes e_j) = \left.\frac{\mathrm{d}}{\mathrm{d}\epsilon}\right|_{\epsilon=0} \det(R)\det(I + \epsilon R^{-1} e_i \otimes e_j).$$

Recalling $\left.\frac{\mathrm{d}}{\mathrm{d}\epsilon}\right|_{\epsilon=0} \det(I + \epsilon A) = \operatorname{tr}(A)$ (all of the off-diagonal entries end up appearing only in ϵ^2 and higher order terms of the determinant, and $\prod_i (1 + \epsilon A_{ii}) = 1 + \epsilon \Sigma_i A_{ii} + O(\epsilon^2)$),

$$\frac{1}{\det(R)}\left[\frac{\mathrm{d} \det(R)}{\mathrm{d} R}\right]_{ij} = \operatorname{tr}(R^{-1} e_i \otimes e_j) = \delta_{km} R^{-1}_{kl} [e_i \otimes e_j]_{lm} = R^{-1}_{kl} [e_i \otimes e_j]_{lk} = R^{-1}_{kl}\delta_{li}\delta_{kj} = [R^{-1}]_{ji}.$$

So $\frac{\mathrm{d}\det(R)}{\mathrm{d}R} = \det(R)[R^{-1}]^T$.