The Computation and Sensitivity of Double Eigenvalues

Ross A. Lippert and Alan Edelman

January 29, 1998

Abstract

In this paper, we address a problem left open by Wilkinson concerning the computation of the shortest (least squares) distance of any matrix A to a matrix $\hat{A}$ with at least one repeated eigenvalue.

This paper presents a method which will identify $\hat{A}$ for a given A based on a geometric interpretation of the level sets of the **eig** function and this method will be based, ultimately, on the examination of the pseudospectra of A.

Additionally, a relation will be given between the conditioning of the pseudospectral method for finding $\hat{A}$ and the conditioning of the problem of finding $\hat{A}$ in general.

1 Introduction

In this paper, we hope to address a problem left open by Wilkinson [sensitivity1] [sensitivity2] concerning the computation of the shortest (least squares) distance of any matrix, A, with distinct eigenvalues to the nearest matrix, $\hat{A}$, with at least one repeated eigenvalue. This distance has been shown to relate to the sensitivity of the eigenvalues of A, in that if the distance to $\hat{A}$ is small, then A will have at least two very ill-conditioned eigenvectors with very ill-conditioned eigenvalues. Although no procedure was given for this computation, Wilkinson reviewed and critiqued several bounds on this distance.

This paper presents a method which will identify the $\hat{A}$ for a given A in most cases, and give appropriate condition numbers on $\hat{A}$ and its repeated eigenvalue. The method proposed also relates the distance from A to $\hat{A}$ to the critical points of the spectral portrait of A (the graph of $z = \sigma(A - (x + iy)I))$ in a very straightforward fashion. In addition we shall find that the conditioning of $\hat{A}$ is directly related to the degeneracy of an associated critical point.

Throughout this paper we will use the Frobenius matrix norm and inner product. We will also assume some passing familiarity with concepts from differential geometry. We will also use some terminology from singularity theory for which [curves and singularities] was an excellent reference.

2 The Level Sets of eig

Let A be a matrix. We define $\mathbf{eig}(A)$ to be some eigenvalue of A. Conversely, we define $M_\lambda = \mathbf{eig}^{-1}(\lambda)$ (where $\lambda = x + iy$) to be the set of all matrices over either R or C with eigenvalue λ. M_λ is a variety which is differentiable almost everywhere.

It is well known that if $\lambda = \mathbf{eig}(A)$ is a simple eigenvalue with right eigenvector x ($Ax = \lambda x$) and left eigenvector y ($y^H A = \lambda y^H$), then variations in λ relate to the variations in A by the formula

$$\dot{\lambda} = \frac{y^H \dot{A} x}{y^H x}.$$

This formula captures the differential behavior of $\mathbf{eig}$ at regular points.

Alternatively, the differential behavior of any scalar (analytic) function $z = f(x)$ can be captured by a gradient and a Riemannian (Hermitian) inner product.

$$\dot{z} = \langle \nabla f, \dot{x} \rangle.$$

Therefore, using inner product $\langle A, B \rangle = \mathrm{tr}(A^H B)$ we may rewrite the perturbation formula for λ as

$$\dot{\lambda} = \langle \frac{yx^H}{x^H y}, \dot{A} \rangle$$

from which we see that the gradient of $\mathbf{eig}$ is actually given by $\frac{yx^H}{x^H y}$.

The gradient of any function it always perpendicular to the iso-surfaces of that function. Thus, $\frac{yx^H}{x^H y}$ is normal to the surface M_λ. If we select y and x such that $||y|| = ||x|| = 1$ then $N = yx^H$ is the unit normal (unique up to sign [or phase]). We may identify $\zeta = x^H y$ as the inverse of the length of the gradient (also up to sign [or phase]).

For points of M_λ with an ill-conditioned eigenvalue, λ, the eigenvectors of λ are close to being perpendicular, thus $\zeta = x^H y$ is very small and the magnitude of the gradient approaches infinity. If A were a matrix with a nontrivial Jordan block for λ then $x^H y$ would vanish and the gradient of $\mathbf{eig}$ would be infinite.

We may alternatively characterize the M_λ surfaces by a different constraint equation with a somewhat better differential behavior. If we let $F(\lambda, A) = \sigma_{min}(A - \lambda I)$, we may define M_λ by the constraint equation $F(\lambda, A) = 0$, λ fixed. The function F has the differential behavior

$$\frac{d}{dt} F(\lambda(t), A(t)) = u^H \dot{A} v - \dot{\lambda} u^H v$$

(note, since a vanishing singular value is not simple, the phases of u and v must be set so as to make the RHS real and positive) Observe that u and v are unit length left and right eigenvectors (respectively), from which we may recover

$$\dot{\lambda} u^H v = \langle uv^H, \dot{A} \rangle.$$

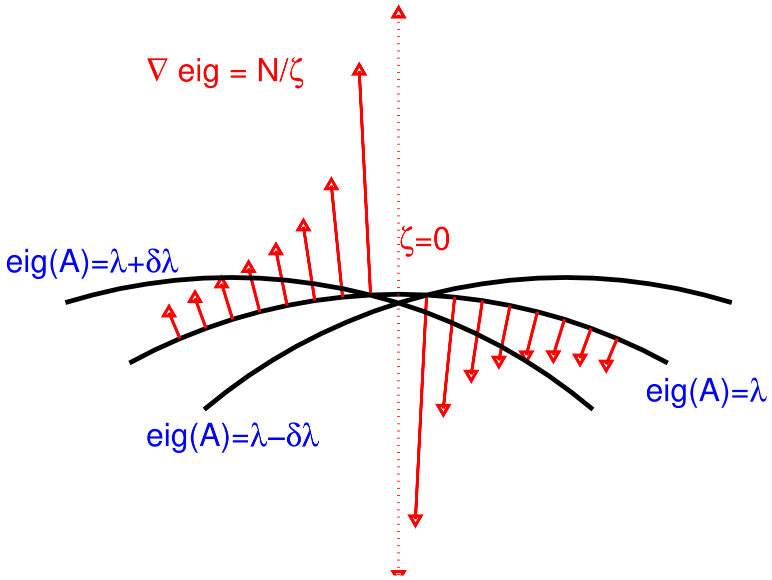

Figure 1: The M_λ parameterized by the constraint $\mathbf{eig}(A) = \lambda$. In this parameterization the gradient points in the direction of increasing λ, and has magnitude inversely proportional to the distance to $M_{\lambda+\delta\lambda}$.

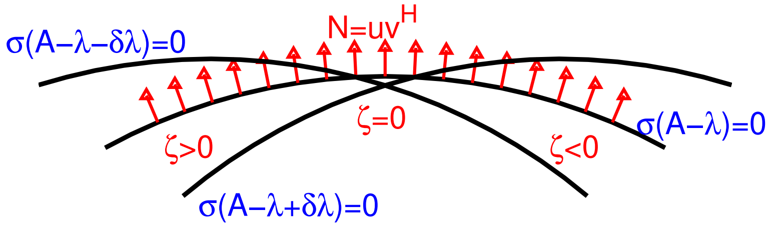

Figure 2: The M_λ can be characterized by singular values, resulting in a natural separation of the N and ζ parts of the gradient of **eig** .

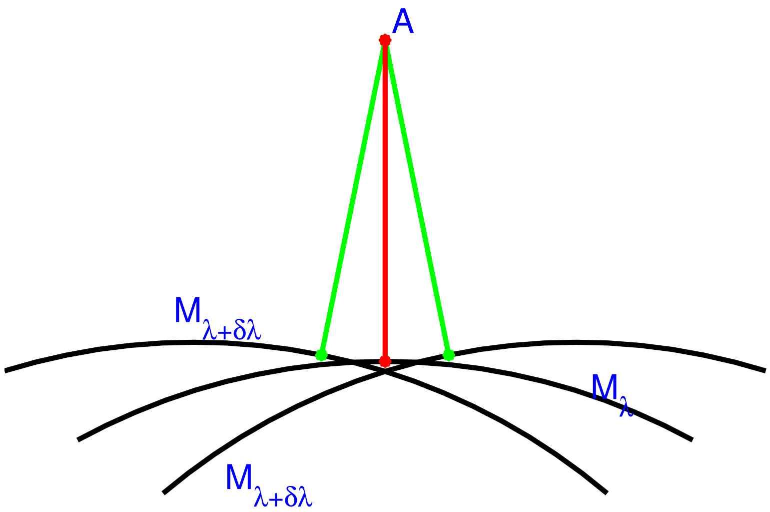

Figure 3: The envelope can be identified by the critical points of the distance of A to M_λ.

3 The Double Eigenvalue Variety

The characterization of the M_λ in terms of $\sigma_{min}(A - \lambda I)$ allows us to analyze the double eigenvalue variety, $\mathcal{D}$, more easily. Recall that $A \in \mathcal{D}$ implies that there exist λ, u, and v such that $Av = \lambda v$, $u^H A = \lambda u^H$, and $u^H v = 0$. Thus, for $A \in \mathcal{D}$ we have

$$F(\lambda, A) = 0$$

for some λ, and for a curve $A(t) \in \mathcal{D}$ with repeated eigenvalue $\lambda(t)$, we have

$$\frac{d}{dt} F(\lambda(t), A(t)) = 0 = u^H \dot{A} v,$$

giving the unit normal $N = uv^H$ of $\mathcal{D}$ at $A(t)$.

From this we observe that $\mathcal{D}$ is the envelope of the family M_λ. Recall that the envelope of a family of surfaces S_t parameterized by t can be defined as a tangent surface E, having the property that for $x \in E$, there exists a t_0 such that $x \in S_{t_0}$ and the tangent space of E at x coincides with the tangent space of S_{t_0} at x.

By using the singular values, we have also recovered an important property of envelopes. $F(\lambda, A)$ is the distance function from A to M_λ. It is the case for all continuously parameterized families of surfaces, that the envelope is given by the the critical points of the distance function, i.e. $\hat{A} \in \mathcal{D}$ when $\frac{d}{dt} F(\lambda(t), A) = 0$ for some A such that $||A - \hat{A}|| = F(\lambda, A)$.

The variety of double eigenvalue matrices, $\mathcal{D}$ can also be characterized as the self-intersections of each M_λ (derogatory matrices) unioned with the set of matrices with non-diagonalizable matrices (degenerate matrices). In the first case, the gradient of **eig** is not unique, in the second, the gradient of **eig** is infinite. However, if we focus on σ_{min} instead of **eig** these difficulties disappear. (In some sense, we have replaced the troubles of ∇**eig** with the troubles of $\nabla \sigma_{min}$, but the latter we find less problematic.)

It is also possible to arrive at these conclusions more algebraically. Let A be a locally differentiable point in $\mathcal{D}$. By taking variations of A within $\mathcal{D}$ we have

$$\delta(u^H Av) = 0 = u^H \delta Av + \lambda(\delta u^H v + u^H \delta v).$$

However, $\delta(u^H v) = 0 = \delta u^H v + u^H \delta v$, thus

$$u^H \delta Av = \langle uv^H, A\rangle = 0,$$

and we again conclude that uv^H is normal to tangents of $\mathcal{D}$ at A.

In some sense, we have arrived at a geometric reasoning for the choice of the "locally fastest" perturbation to a matrix A near $\mathcal{D}$ which would move A to $\mathcal{D}$ nearly optimally. From the preceding results, we see that if $\hat{A} \in \mathcal{D}$ then the normal line $\hat{A} + suv^H$ passes through $\mathcal{D}$ along a perpendicular, thus the perturbation $E = suv^H$ can be thought of as the perturbation which maximizes the change in the distance to $\mathcal{D}$ for small s.

Alternatively, if one suspects that two eigenvalues, λ_1, λ_2, of A can be coalesced with a small perturbation E, Wilkinson suggested using a perturbation of the form $E = s(y_1 + y_2)(x_1+x_2)^H$ where the y_i and the x_i are the left and right eigenvectors of λ_i and s is adjusted so that $A + E \in \mathcal{D}$ (section 9 of [Wilkinson1]).

If one is close to $\mathcal{D}$ then one expects $(x_1 + x_2)/2$ will be close to v and $(y_1 + y_2)/2$ will be close to u. The "locally fastest" perturbation is then expected to be close to the actual normal of the variety at the actual minimizing point. We will elaborate on this heuristic later.

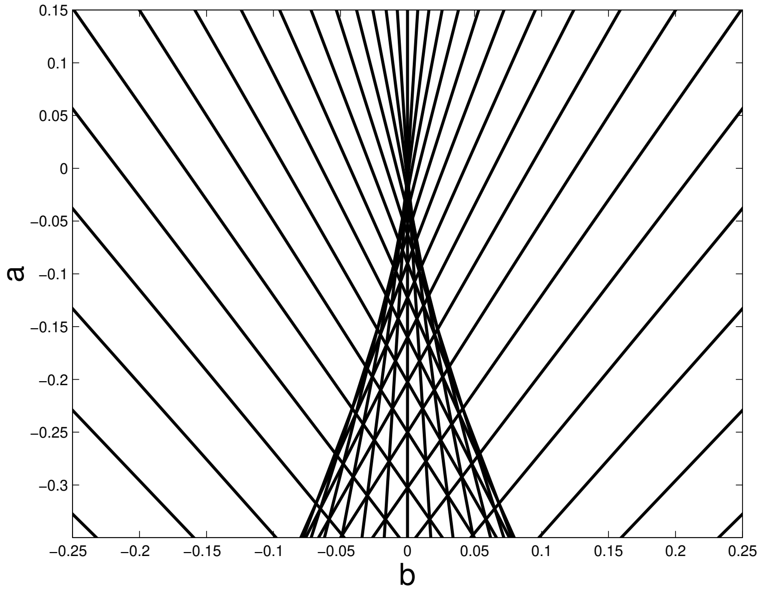

Figure 4: A graph in the a-b plane of the family of polynomials of the form $p(x) = x^3 + ax + b$ where $p(\lambda) = 0$.

4 An Illustration by Characteristic Polynomials

For illustrative purposes, we may consider the cubic polynomials of the form $p(x) = x^3 + ax + b$. If we select the family of cubics of this form by a root λ, then polynomials in the root λ can be represented by a linear relation between a and b of the form

$$\lambda^3 + a\lambda + b = 0.$$

A collection of these lines is shown in the figure.

The lines trace out a cusp-like shape. By varying the constraint equation

$$3\lambda^2\dot{\lambda} + a\dot{\lambda} + \dot{a}\lambda + b = 0,$$

we obtain the gradient of λ as a function of a and b,

$$\dot{\lambda} = \frac{-(\lambda, 1) \cdot (\dot{a}, \dot{b})}{3\lambda^2 + a}.$$

We see that we may define a unit normal $N = \frac{\lambda x + 1}{\sqrt{\lambda^2+1}}$ and set the inverse magnitude of the gradient to $\zeta = \frac{3\lambda^2+a}{\sqrt{\lambda^2+1}}$.

The envelope of the lines of this single root family of polynomials is given by $\zeta = 0$, or $a = -3\lambda^2$ (and thus $b = 2\lambda^3$ by the constraint). One can check that this gives a curve through polynomial space of polynomials with double root λ, plotted in the figure (along with a few of its the normal lines). The astute reader may observe that the actual cusp point of this curve is the polynomial $p(x) = x^3$, the triple root polynomial.

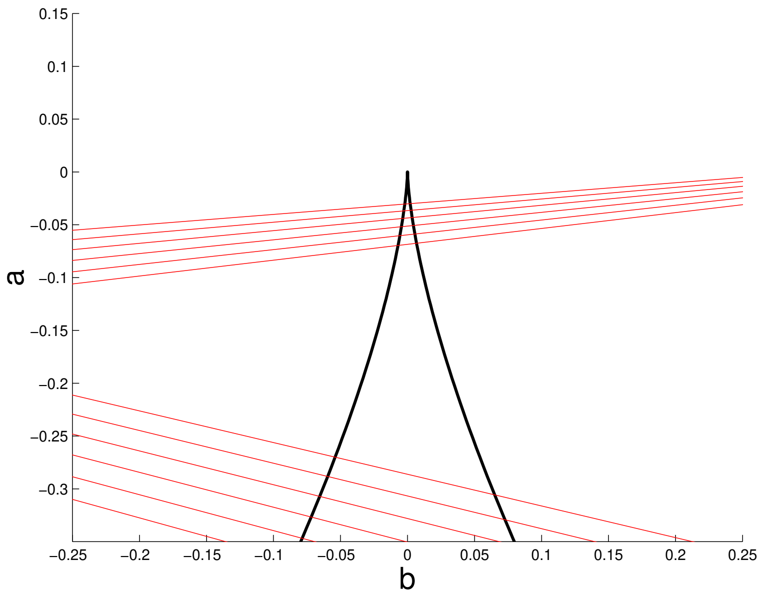

Figure 5: A graph in the a-b plane of the envelope of the single root family of polynomials.

The connection between the double eigenvalue variety and the double root variety can be made with the 1-dimensional family of matrices given by

$$M(\lambda) = \begin{pmatrix} \lambda & \sqrt{1-2\lambda^2} & \lambda \\ 0 & -2\lambda & \sqrt{1-2lambda^2} \\ 0 & 0 & \lambda \end{pmatrix}.$$

One can show that $poly(M(\lambda);x) = x^3 - 3\lambda^2 x + 2\lambda^3$. Also, it is clear that $uv^H = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}$. One can verify that $poly(M(\lambda)+suv^H;x) - poly(M(\lambda);x) = s(\lambda x+1)$, parallel to the normals of the double root polynomials.

We will often use characteristic polynomials of specially selected matrices as illustrative examples of the geometry of $\mathcal{D}$ in what follows.

5 A Primal-Dual Attempt to Solve the Nearest Double Eigenvalue Problem

Edelman and Elmroth introduced an algorithm to find the nearest $\hat{A} \in \mathcal{D}$ to a given matrix A. The algorithm is the simultaneous solution of two dual problems which can be stated as follows:

Problem 1:
Given a matrix A and an approximation to the perturbation bringing A to $\mathcal{D}$, find the repeated eigenvalue of $\hat{A}$.

Problem 2:
Given a matrix A and an approximation of the repeated eigenvalue of $\hat{A}$, find the smallest perturbation bringing A to $\mathcal{D}$.

For Problem 1, they used a clustering algorithm on the eigenvalues of $A - suv^H$, where suv^H is the approximation to the normal bringing A to $\mathcal{D}$ (s a scalar and uv^H an approximate normal). For Problem 2, they approximated u and v with the left and right singular vectors of $A - \lambda I$ for the given approximate repeated eigenvalue λ.

This resulted in an extremely compact matlab routine, which reads roughly as

```
lam = guess;
while (1)
[u,s,v] = svd(A-lam*I);
e = eig(A - u(:,n)*s(n,n)*v(:,n)');
lam = cluster_routine(e);
end
```

where the cluster routine took the mean of two nearest elements of `e`.

The resulting algorithm converges linearly much of the time, but it will sometimes go into cycles which are empirically not convergent. At present, we do not have a detailed analysis of the convergence of the primal-dual algorithm.

The following figures show the primal-dual algorithm both succeeding and failing for two different matrices. We have plotted the new value of the variable `e` against its value in the previous iteration.

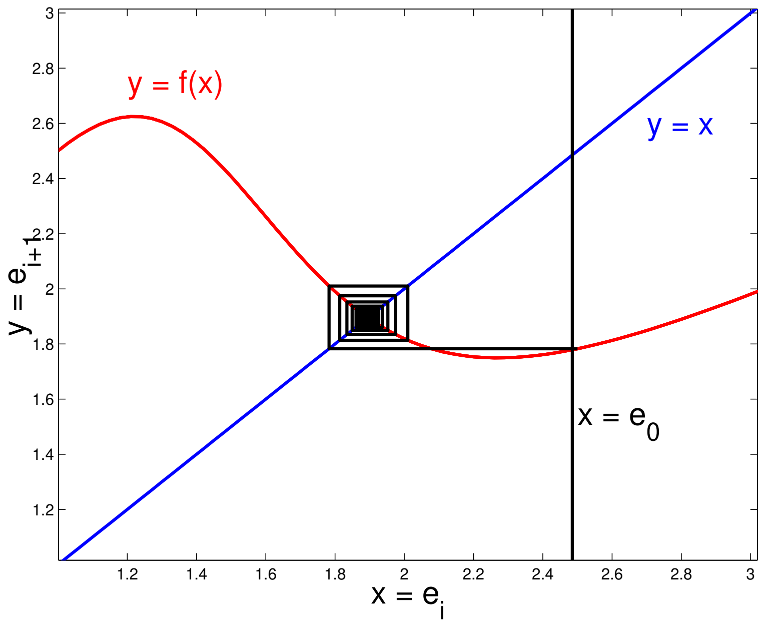

Figure 6: Here we see the primal-dual algorithm converging.

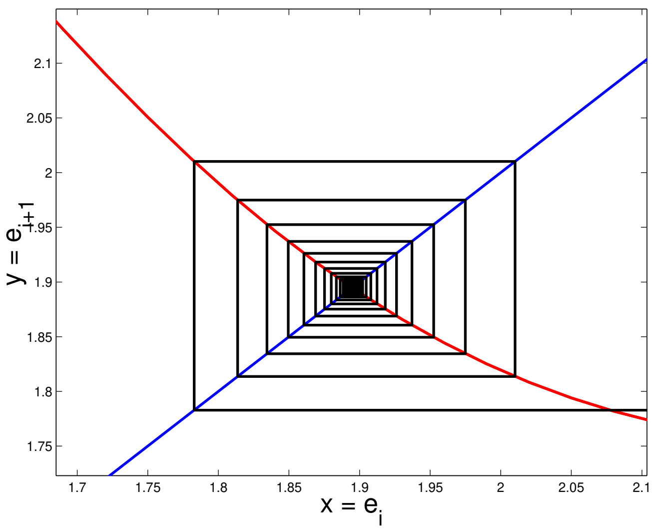

Figure 7: A blow-up of the convergent case.

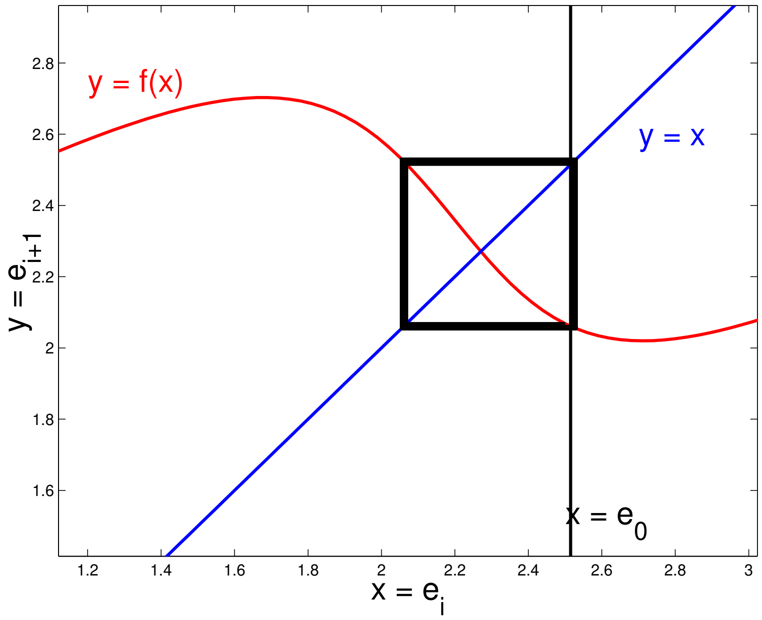

Figure 8: Here we see the primal dual algorithm converging to some sort of limit-cycle.

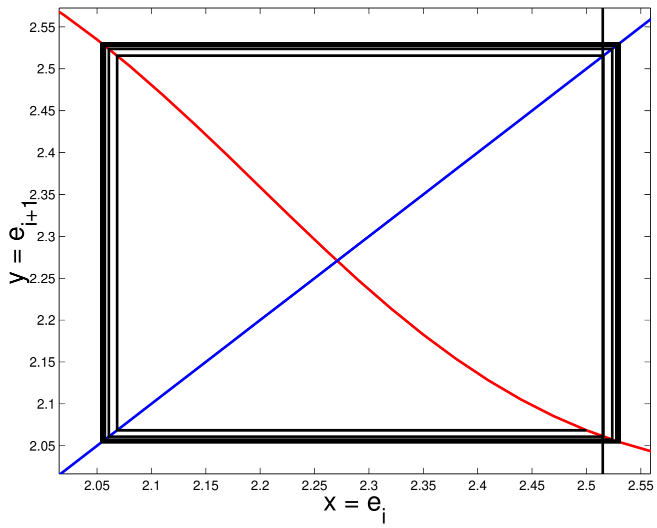

Figure 9: A blow-up of the limit-cycle case.

6 A Pseudo-Spectral Interpretation

We have identified that the normal space of a matrix $\hat{A} \in \mathcal{D}$ is given by $\hat{A} + suv^H$, where u, v are left and right eigenvectors of $\hat{A}$. Thus if A lies in the normal space if its nearest matrix $\hat{A} \in \mathcal{D}$, must satisfy $A = \hat{A} + suv'$. There is, however, no obvious way to compute u and v without knowing $\hat{A}$ in advance. In this section, we present a method for this based on the pseudospectra of A.

Let $\sigma(x,y) = \sigma_{min}(A - (x+iy)I)$. In the case where $\sigma(x,y)$ is smooth, we have the following theorem about the derivatives of $\sigma(x,y)$ due to Sun[].

Theorem 6.1 *Let $U(x,y)\Sigma(x,y)V^H(x,y) = A - (x+iy)I$, where $\sigma_n(x_0,y_0)$ is simple, and let $\tilde{U} = (\, u_1 \quad u_2 \quad \ldots \quad u_{n-1}\,)$, $\tilde{V} = (\, v_1 \quad v_2 \quad \ldots \quad v_{n-1}\,)$, and $\tilde{\Sigma} = diag(\sigma_1, \ldots, \sigma_{n-1})$.*

$$
\begin{aligned}
\frac{\partial \sigma}{\partial x}(x_0,y_0) &= -Re\{u_n^H v_n\} \\
\frac{\partial \sigma}{\partial y}(x_0,y_0) &= Im\{u_n^H v_n\}
\end{aligned}
$$

$$
\begin{aligned}
\frac{\partial^2 \sigma}{\partial x^2}(x_0,y_0) &= Re\{r^H \Phi r + l^H \Phi l + 2l^H \Psi r\} + Im\{u_n^H v_n\}^2/\sigma_n \\
\frac{\partial^2 \sigma}{\partial y \partial x}(x_0,y_0) &= -Im\{2l^H \Psi r\} + Im\{u_n^H v_n\} Re\{u_n^H v_n\}/\sigma_n \\
\frac{\partial^2 \sigma}{\partial y^2}(x_0,y_0) &= Re\{r^H \Phi r + l^H \Phi l - 2l^H \Psi r\} + Re\{u_n^H v_n\}^2/\sigma_n
\end{aligned}
$$

where $\Phi = \sigma_n(\sigma_n^2 I - \tilde{\Sigma}^2)^{-1}$, $\Psi = \tilde{\Sigma}(\sigma_n^2 I - \tilde{\Sigma}^2)^{-1}$, $r = \tilde{U}' v_n$, and $l = \tilde{V}' u_n$.

Sun's theorem implies that if $\sigma_x = \sigma_y = 0$, then $u^H v = 0$. Thus the matrix $\hat{A} = A - \sigma(x_0,y_0)uv^H \in M_{x_0+iy_0}$ is actually a point on the envelope of the family of M_λ, i.e. $\hat{A} \in \mathcal{D}$. The stationary points of $\sigma(x,y)$ correspond to normal lines off of $\mathcal{D}$. Thus, an examination of the pseudospectra of A can be used to determine the distance from A to $\hat{A} \in \mathcal{D}$.

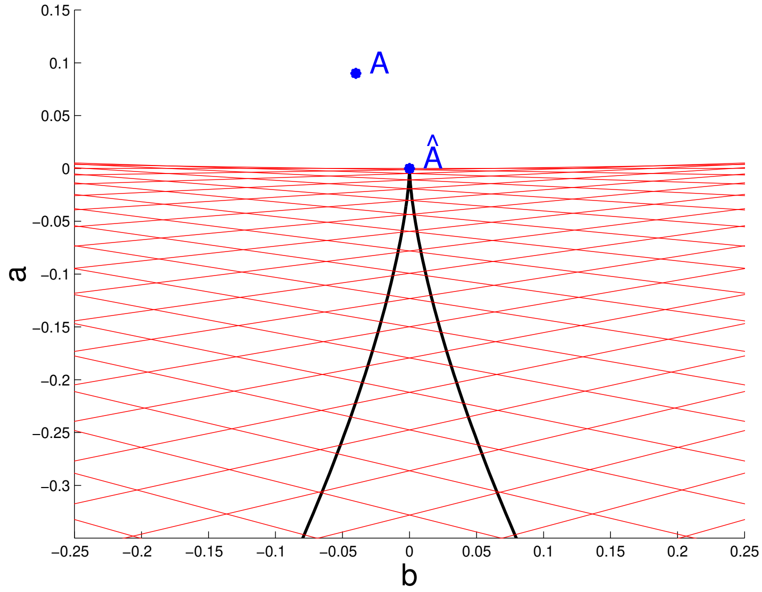

Figure 10: A possible problem finding a normal from $\hat{A}$ to A illustrated in characteristic polynomial space.

7 Technical Lemmae

In the previous section, we have showed that with some possible assumptions of generic conditions on A the nearest matrix, $\hat{A} \in \mathcal{D}$, to A is given by $A - suv^H$ where $s = \sigma_{min}(A - \lambda I)$ is stationary in λ, u and v are the associated singular vectors and λ is a repeated eigenvalue of $\hat{A}$ for which u and v are right and left eigenvectors.

If one were a bit cavalier, one might assume that the nearest element of $\mathcal{D}$ was always some degenerate matrix with a J_2 structure and no further nongeneric properties. If this were the case, then $\mathcal{D}$ would be locally differentiable about $\hat{A}$, consequently there would be a unique normal perturbation, and that perturbation would be of the form uv^H where $u^H v = 0$. Thus, the $\hat{A}$ would correspond to a critical point of the pseudospectra of $A - \lambda I$ as argued in the previous section.

However, one might consider the ways in which this generic assumptions could be violated. If $\hat{A}$ were derogatory or had degeneracy less generic than J_2, then there would be multiple normal directions off of $\hat{A}$, and $A - \hat{A}$ might not be of the assumed form. One particular example of where one might really be concerned about this is in the case where $\hat{A}$ actually has a J_3 structure and A might sit above all of rank 1 normals of $\hat{A}$ as shown in the figure (since a matrix with J_3 structure is actually a cusp point of D, there is no reason to believe that it has a continuous local normal field).

Clearly, one needs to have some lemmae to assure understand the conditions under which such an event will or will not happen in order to make generalizations which are everywhere valid. The purpose of this section is to clarify these possibilities.

It will be useful to take a partial canonical decomposition of $\hat{A}$ in the repeated eigenvalue.

That is, let

$$\hat{A} = \begin{pmatrix} E & X \end{pmatrix} \begin{pmatrix} J(\lambda) & 0 \\ 0 & M \end{pmatrix} \begin{pmatrix} \mathcal{E} & Y \end{pmatrix}^H,$$

where $\begin{pmatrix} \mathcal{E} & Y \end{pmatrix}^H \begin{pmatrix} E & X \end{pmatrix} = I$ and $J(\lambda)$ is the canonical matrix of the Jordan structure of $\hat{A}$ at λ.

Lemma 7.1 *Let $\hat{A}$ be the nearest matrix in $\mathcal{D}$ to A. Then*

$$N = A - \hat{A} = \mathcal{E} W E^H,$$

for some W.

Proof: We also decompose $N = A - \hat{A}$ according to the following (slightly twisted) decomposition,

$$N = \begin{pmatrix} \mathcal{E} & X \end{pmatrix} \begin{pmatrix} N_{11} & N_{12} \\ N_{21} & N_{22} \end{pmatrix} \begin{pmatrix} E & Y \end{pmatrix}^H,$$

one may check that $\begin{pmatrix} E & Y \end{pmatrix}^H \begin{pmatrix} \mathcal{E} & X \end{pmatrix} = I$, and that the displacements $\mathcal{E} N_{11} E^H$, $X N_{12} E^H$, $\mathcal{E} N_{21} Y^H$, and $X N_{22} Y^H$ are mutually orthogonal.

Since $\hat{A} + X N_{22} Y^H$ has the same canonical structure in λ as $\hat{A}$, it is clear that $N_{22} = 0$ for $\hat{A}$ to be the minimizer. Similarly, by observing that $(\hat{A} + \mathcal{E} N_{21} Y^H) E = EJ$ and $\mathcal{E}^H(\hat{A} + X N_{12} E^H) = J \mathcal{E}^H$, we see that N_{12} and N_{21} must vanish. $\square$

For the next lemma, we consider the case where J the Jordan structure of $\hat{A}$ at λ is a product of the form $J = \begin{pmatrix} J_1 & 0 \\ 0 & J_2 \end{pmatrix}$ and the eigenvectors are likewise partitioned into $\mathcal{E} = \begin{pmatrix} \mathcal{E}_1 & \mathcal{E}_2 \end{pmatrix}$ and $E = \begin{pmatrix} E_1 & E_2 \end{pmatrix}$. Now we decompose N as

$$N = \begin{pmatrix} \mathcal{E}_1 & \mathcal{E}_2 \end{pmatrix} \begin{pmatrix} N_{11} & N_{12} \\ N_{21} & N_{22} \end{pmatrix} \begin{pmatrix} E_1 & E_2 \end{pmatrix}^H.$$

Lemma 7.2 *Suppose that the Jordan structure of $\hat{A}$ is a product. Then N is of the form,*

$$N = \mathcal{E}_1 N_{11} E_1^H + \mathcal{E}_2 N_{22} E_2^H.$$

Proof: Observe that $\hat{A} + E_2 W \mathcal{E}_1^H$ has the Jordan structure of $\begin{pmatrix} J_1 & W \\ 0 & J_2 \end{pmatrix}$ which, while being possibly different from J, still has a repeated eigenvalue λ.

$\hat{A} + E_2 W \mathcal{E}_1$ is closer to A than $\hat{A}$ if $\langle E_2 W \mathcal{E}_1, N \rangle = \langle W, N_{12} \rangle$ is greater than zero for some arbitrarily small W. Thus, $N_{12} = 0$.

By a similar argument, $N_{21} = 0$. $\square$

Finally, we can, in a sense, eliminate the derogatory $\hat{A}$'s by the following lemma.

Lemma 7.3 *If $\hat{A}$ has a product canonical structure, J, then one of the irreducible blocks of J, say J_1 is degenerate, and*

$$N = \mathcal{E}_1 N_{11} E_1^H.$$

Proof: Suppose that all of the irreducible blocks of J were simple, i.e. $J = \lambda I$. Then, from the previous lemma, $N = \epsilon_1 n e_1^H$ (where e_1 and ϵ_1 are vectors and n is a scalar).

$\hat{A} + \mu I$ has a repeated eigenvalue of $lambda + \mu$.

$\hat{A} + \mu I$ will be closer to A if $\langle \mu I, N \rangle = \mathrm{Re}\{\bar{\mu} n\}$ is greater than zero for some arbitrarily small μ. Thus $n = 0$ and $\hat{A} = A$ contradicting hypothesis.

Combining this with the result of the previous lemma gives the the conclusion. □

We summarize with the following corollary.

Corollary 7.4 *The nearest matrix $\hat{A} \in \mathcal{D}$ to A is reached by a perturbation of the form*

$$N = \mathcal{E} p(J^t) E^H,$$

where $p(x)$ is a $k-1$-degree polynomial with vanishing constant term, and E, $\mathcal{E}$ are the right and (corresponding dual) left generalized eigenvectors of a canonical $k \times k$ Jordan block, J, in the Jordan structure of $\hat{A}$.

Proof: Prior lemmae have shown that N is of the form $N = \mathcal{E} W E^H$. The similarity transformation $e^{sX} \hat{A} e^{-sX} \in \mathcal{D}$ is closer to A for some arbitrarily small s whenever the commutator, $X\hat{A} - \hat{A}X$ has positive inner product with N.

Thus $\langle [X, \hat{A}], N \rangle = \langle X, [N, \hat{A}^H \rangle$ must vanish for all X. Giving $N\hat{A}^H - \hat{A}^H N = 0$ which implies that $WJ^t = J^t W$, and thus $W = p(J^t)$. □

We now need only really consider the case where the Jordan structure of $\hat{A}$ is a single canonical Jordan block.

Lemma 7.5 *Let the Jordan block J of $\hat{A}$ be of size greater than 2. The polynomial $p(x)$ in*

$$N = \mathcal{E} p(J^t) E^H,$$

must be of the form $p(x) = ax^{k-1} + bx^{k-2}$.

Proof: Suppose that the coefficient, c, of x^j ($j < k-2$) is nonvanishing. Then we consider the matrix $\hat{A} + EM\mathcal{E}^H$, where $M_{kj} = s$ ($M_{pq} = 0$ elsewhere).

$\hat{A} + EM\mathcal{E}^H$, while not necessarily having the same canonical structure as J still has λ repeated at least twice, since

$$(\hat{A} + EM\mathcal{E}^H) \begin{pmatrix} e_1 & e_2 \end{pmatrix} = \begin{pmatrix} e_1 & e_2 \end{pmatrix} \begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix}.$$

However, $\hat{A} + EM\mathcal{E}^H$ is closer to A so long as $\langle M, p(J^t) \rangle = sc$ does not vanish for arbitrarily small s. Thus we must have $c = 0$. □

Thus, for all A not in $\mathcal{D}$, $N = A - \hat{A}$ is of the form that one generically expects from either a J_2 or a J_3 canonical structure. In what follows, we shall see that if we allow $\hat{A}$ to take values in C then the $A - \hat{A}$ will be rank 1.

Theorem 7.6 *Let $\hat{A} \in \mathcal{D}$ be the nearest matrix with a repeated eigenvalue to the given matrix A with distinct eigenvalues. Then*

$$N = A - \hat{A} = suv^H,$$

where u and v are the left and right eigenvectors of $\hat{A}$ and s is a scalar.

Proof: If the Jordan structure of the repeated eigenvalue of $\hat{A}$ is a canonical J_2 block, then $\mathcal{D}$ is locally differentiable about $\hat{A}$, thus A must lie on a normal to $\hat{A}$.

Assuming the canonical structure of $\hat{A}$ is $J = J_k(\lambda)$, we have from the previous lemmae, we have that

$$N = \mathcal{E}(a(J^t)^{k-1} + b(J^t)^{k-2})E^H,$$

$b \neq 0$.

Consider the matrix $\hat{A}+E(y(J^t)^{k-1}+x(J^t)^{k-2})\mathcal{E}^H$, where $x = k\mu^{k-1}$ and $y = -(k-1)\mu^k$. It is easy to show that $J + (y(J^t)^{k-1} + x(J^t)^{k-2})$ has a repeated (J_2) eigenvalue $\lambda + \mu$.

$\hat{A} + E(y(J^t)^{k-1} + x(J^t)^{k-2})\mathcal{E}^H$ will be closer to A then $\hat{A}$ if for some arbitrarily small μ $\langle E(y(J^t)^{k-1} + x(J^t)^{k-2})\mathcal{E}^H, N\rangle$ is greater than zero.

$\langle E(y(J^t)^{k-1}+x(J^t)^{k-2})\mathcal{E}^H, N\rangle = \mathrm{Re}\{\bar{x}b+\bar{y}a\}$. Substituting for x and y we have $\mathrm{Re}\{\bar{x}b+\bar{y}a\} = \mathrm{Re}\{\bar{\mu}^{k-1}(kb-(k-1)\bar{\mu}a\}$. If we take $|\mu| < \frac{b}{a}$ and let select the phase of μ such that $\mu^{k-1}b$ is positive real, then $\langle E(y(J^t)^{k-1} + x(J^t)^{k-2})\mathcal{E}^H, N\rangle > 0$.

Thus, we have $\hat{A} + E(y(J^t)^{k-1} + x(J^t)^{k-2})\mathcal{E}^H$ closer to A than $\hat{A}$.

We must conclude that $b = 0$ and therefore that N is rank 1 of the form asserted. □

The reader will note that it is not always possible to find such x and y if one is constrained to real matrices, but one can always do so over complex numbers. In fact, the illustration at the beginning of this section demonstrates this shortcoming of the reals. If this illustration had been drawn with complex a and b axes, then it would have been apparent that A was actually closer to a complex $\hat{A}$.

We have now established rigorously that any matrix A with distinct eigenvalues must lie along a rank 1 normal from the nearest $\hat{A} \in \mathcal{D}$. Thus $\hat{A}$ can be determined by examining the singular values of $A - \lambda I$.

What remains are now some lemmae concerning the singular values themselves so that it can be determined what sort of behavior one must look for in the singular values in order to locate $\hat{A}$. The previous section claims that the critical points of σ_{min} are the points of interest in determining $\hat{A}$. However, one may wonder what happens when σ_{min} is not simple (and hence not strictly differentiable), and whether or not one must examine the critical behavior of other singular values.

The first concerns the local behavior of $\sigma_{min}(A - (x + iy)I)$ when σ_{min} is not simple.

Lemma 7.7 *Let $x_0 + iy_0$ be a point the complex plane where $s = \sigma_{min}(A - (x_0 + iy_0)I)$, and*

$$(A - (x_0 + iy_0)I)V = sU$$
$$U^H(A - (x_0 + iy_0)I) = sV^H.$$

if $\exists x$ such that $x^H U^H V x = 0$. $\exists u, v, s$ such that $\hat{A} = A - suv^H \in \mathcal{D}$

Proof: Given x, we may take $||x|| = 1$. Let $u = Ux$ and $v = Vx$. We select an orthogonal matrix Q with first column equal to x. We may then replace U with UQ and V with VQ in the svd of $(A - (x_0 + iy_0)I$, the first column of U being u and the first column of V being v, and $u^H v = 0$ according to hypothesis. Thus $A - suv^H \in \mathcal{D}$. □

We see that it is possible to find $\hat{A}$ from nonsimple σ_{min}. Thus, critical points of the spectral portrait, are expect to the more generic means of finding $\hat{A}$, this theorem tells us that the crossing points of σ_{min} with other singular values are also possible sources of $\mathcal{D}$ normals.

Lastly, although we have not found a counterexample, we have no proof that σ_{min} is the only singular value to examine. It may be that $s_0 = \sigma_n(A - (x_0 + iy_0)I)$ and $s_1 = \sigma_{n-1}(A - (x_1 + iy_1)I)$ are both stationary with $s_1 < s_0$ for some A. This would mean that all singular values must be inspected for crossing points and critical points, not just σ_{min}. However, we have attempted to find such matrices and have failed, so, for now, this is an *empirical lemma.*

8 Wilkinson Revisited

A normal perturbation off of a generic point $\hat{A} \in \mathcal{D}$ (i.e. has exactly one J_2 Jordan block) results in a diagonalizable matrix $A = \hat{A} + suv^H$. The J_2 block that was part of the structure of $\hat{A}$ splits into two J_1 blocks for $|s| > 0$. Wilkinson believed that the average of the two eigenvectors and eigenvalues of A would be good approximations to the eigenvalue and eigenvector of $\hat{A}$, though he was not able to present a rigorous proof of this. In this section we shall study the resulting eigenspaces of A in terms of the double eigenvalue decomposition and derive the Puiseux series in terms of the variable $\sqrt{s}$.

We begin with the decomposition of a matrix A on the normal of a matrix with a double eigenvalue λ, with right and left eigenvectors v and u:

$$A = suv^H + \lambda I + \tilde{U}\tilde{\Sigma}\tilde{V}^H.$$

We look for right eigenvectors of the form $x(s) = v + \tilde{V}q(s) = v + \tilde{V}(\sqrt{s}q_1 + sq_2 + \ldots)$ and eigenvalues of the form $\lambda + \eta(s)$ (the form of $\eta(s)$ will be given later). Substituting into $A = (\lambda + \eta(s))x(s)$ we have:

$$\begin{aligned} A(v + \tilde{V}q(s)) &= (\lambda + \eta(s))(v + \tilde{V}q(s)) \\ su + \tilde{U}\tilde{\Sigma}q(s) + \lambda(v + \tilde{V}q(s)) &= (\lambda + \eta(s))(v + \tilde{V}q(s)) \\ su + \tilde{U}\tilde{\Sigma}q(s) &= \eta(s)(v + \tilde{V}q(s)). \end{aligned}$$

Multiplying both sides by u^H, we have

$$s = \eta(s)l^H q(s),$$

which we use to substitute for $\eta(s)$ and have,

$$su + \tilde{U}\tilde{\Sigma}q(s) = \frac{s}{l^H q(s)}(v + \tilde{V}q(s)),$$

which is automatically satisfied in the u direction. For the $\tilde{U}$ directions we may multiply both sides by $\tilde{\Sigma}^{-1}\tilde{U}^H$ giving

$$(l^H q(s))q(s) = s\tilde{\Sigma}^{-1}\tilde{U}^H(v + \tilde{V}q(s))$$

Equating the first two terms of equal power in $\sqrt{s}$ we have

$$\begin{aligned} (l^H q_1)q_1 &= \tilde{\Sigma}^{-1}r, \\ (l^H q_2)q_1 + (l^H q_1)q_2 &= \tilde{\Sigma}^{-1}\tilde{U}^H\tilde{V}q_1 \end{aligned}$$

which are solved by

$$q_1 = \frac{\tilde{\Sigma}^{-1}r}{\sqrt{l^H\tilde{\Sigma}^{-1}r}}, q_2 = \frac{1}{\sqrt{l^H\tilde{\Sigma}^{-1}r}}\tilde{\Sigma}^{-1}\tilde{U}^H\tilde{V}q_1 - \frac{l^H\tilde{\Sigma}^{-1}\tilde{U}^H\tilde{V}q_1}{2l^H\tilde{\Sigma}^{-1}r}q_1.$$

As expected, there are two solutions x_1 and x_2, corresponding to the sign taken for $\sqrt{l^H\tilde{\Sigma}^{-1}r}$. One may observe that $\tilde{V}\tilde{\Sigma}^{-1}\tilde{U}^H = \hat{A}^{-\perp}$ is the least-squares inverse of $\hat{A}$, and thus

$\tilde{V}\tilde{\Sigma}^{-1}r = w$ is the generalized right eigenvector of $\hat{A}$ perpendicular to v. Incorporating those facts into the formulae we have

$$\tilde{V}q_1 = \frac{w}{\sqrt{u^H w}}, \tilde{V}q_2 = \frac{1}{u^H w}\hat{A}^{-\perp}w - \frac{u^H\hat{A}^{-\perp}w}{2(u^H w)^2}w.$$

If we take the Wilkinson heuristic and average x_1 and x_2 we see that $(x_1 + x_2)/2 = v + s\tilde{V}q_2 + \ldots$, with the $\sqrt{s}$ terms canceling out. This also tells us that the average will be a good approximation so long as $s << \frac{1}{||q_2||}$. We may also note that $||q_2|| = \infty$ when $u^H w = 0$, i.e. when the generalized eigenvector is also perpendicular to the left eigenvector, implying that $\hat{A}$ has a triple eigenvalue.

The formulae for y_1 and y_2 are similar (taking transposes and replacing V's with U's and so on). With formulae for both right and left eigenvectors, we can form the expansion for the cosines of the angles between the right and left eigenvectors. They are (to order s),

$$\cos(x_i, y_i) = |\sqrt{s}(\sqrt{v^H t} + \sqrt{w^H u}) + s(\frac{w^H t}{\sqrt{w^H u}\sqrt{v^H t}} + \frac{w^H(\hat{A}^{-\perp})^H u}{2w^H u} + \frac{v^H(\hat{A}^{-\perp})^H t}{2v^H t}) + \ldots|,$$

$$\cos(x_1 + x_2, y_1 + y_2) = s|(\frac{w^H(\hat{A}^{-\perp})^H u}{2w^H u} + \frac{v^H(\hat{A}^{-\perp})^H t}{2v^H t}) + \ldots|,$$

where the generalized left eigenvector is denoted by t.

We see that while the cosine of the average grows proportionally with s while the cosine of the individual vectors grow as $\sqrt{s}$ to lowest order. Thus we expect the former cosine to be a much better indicator of distance to $\mathcal{D}$

9 Conditioning of the Critical Points of $\sigma(x,y)$

We now consider the sensitivity of the critical points of $\sigma(x,y)$ to small perturbations in A. To first order, a perturbation, δA, in A will cause perturbation in the gradient of $\sigma(x,y) = (\sigma_x, \sigma_y)$ at the point x_0, y_0, which we may call $(\delta\sigma_x, \delta\sigma_y)$. By inverting the Hessian (given by Sun's formulae above) on this perturbation, one then has the first order correction to x_0, y_0, which we may call $\delta x_0, \delta y_0$.

Lemma 9.1 *Given x_0, y_0 is a stationary point of $\sigma(x,y)$ and Φ, Ψ, r, and l be defined as above. Then let $p = \tilde{U}^H \delta A v_n$ and $q = \tilde{V}^H \delta A^H u_n$. then*

$$\delta\sigma_x(x_0,y_0) = Re\{r^H \Phi p + r^H \Psi q + p^H \Psi l + q^H \Phi l\}$$
$$\delta\sigma_y(x_0,y_0) = Im\{r^H \Phi p + r^H \Psi q + p^H \Psi l + q^H \Phi l\}$$

Proof:

From the Sun's theorem, we have

$$\sigma_x - i\sigma_y = -u_n^H v_n$$

and thus

$$\delta\sigma_x - i\delta\sigma_y = -\delta u_n^H v_n - u_n^H \delta v_n.$$

Since $u_n^H v_n = 0$, $v_n = \tilde{U}\tilde{U}^H v_n = \tilde{U} r$ and $u_n = \tilde{V}\tilde{V}^H u_n = \tilde{V} l$. Substituting, we have,

$$\delta\sigma_x - i\delta\sigma_y = -(\tilde{U}^H \delta u_n)^H r - l^H(\tilde{V}^H \delta v_n).$$

Varying $U\Sigma V^H = A - (x_0 + iy_0)I$ we have

$$\delta U \Sigma V^H + U \delta\Sigma V^H + U\Sigma\delta V^H = \delta A$$

and thus

$$(U^H \delta U)\Sigma + \delta\Sigma + \Sigma(V^H \delta V)^H = U^H \delta A V.$$

Equating the bottom row and the rightmost column, we have

$$\Sigma\delta U^H u_n + \begin{pmatrix} 0 \\ 0 \\ \dots \\ 0 \\ \delta\sigma_n \end{pmatrix} + (V^H \delta v_n)\sigma_n = \begin{pmatrix} q \\ v_n^H \delta A^H u_n \end{pmatrix}$$

$$\Sigma\delta V^H v_n + \begin{pmatrix} 0 \\ 0 \\ \dots \\ 0 \\ \delta\sigma_n \end{pmatrix} + U^H \delta u_n \sigma_n = \begin{pmatrix} p \\ u_n^H \delta A v_n \end{pmatrix}.$$

Since the matrices $U^H \delta U$ and $V^H \delta V$ are anti-Hermitian, $\delta U^H u_n = -U^H \delta u_n$, and $\delta V^H v_n = -V^H \delta v_n$. We substitute and take the first $n-1$ rows of both equations:

$$
\begin{aligned}
(\tilde{V}^H \delta v_n)\sigma_n - \Sigma(\tilde{U}^H \delta u_n) &= q \\
-\Sigma(\tilde{V}^H \delta v_n) + (\tilde{U}^H \delta u_n)\sigma_n &= p
\end{aligned}
$$

One may then solve the system of equations given for

$$
\begin{aligned}
\tilde{V}^H \delta v_n &= \Phi q + \Psi p \\
\tilde{U}^H \delta u_n &= \Psi q + \Phi p
\end{aligned}
$$

and then substitute into

$$\delta\sigma_x - i\delta\sigma_y = -(\tilde{U}^H \delta u_n)^H r - l^H(\tilde{V}^H \delta v_n)$$

to obtain the result. $\square$

With this lemma, finding κ then becomes a question of finding the largest singular value of the $2 \times (4n-4)$ real matrix

$$\begin{pmatrix} \sigma_{xx} & \sigma_{xy} \\ \sigma_{yx} & \sigma_{yy} \end{pmatrix}^{-1} \begin{pmatrix} \delta\sigma_x \\ \delta\sigma_y \end{pmatrix}$$

where $\delta\sigma_x$ and $\delta\sigma_y$ are the real linear functions of the real and imaginary parts of p and q defined by the lemma.

10 Conditioning of $\hat{A}$

In the previous section we found that finding the critical points of the spectral portrait of A can be ill-conditioned when the hessian of $\sigma(x_0, y_0)$ is singular. This occurs when

$$\mathrm{Re}\{r^H \Phi r + l^H \Phi l\}^2 - \mathrm{Re}\{2l^H \Psi r\}^2 - \mathrm{Im}\{2l^H \Psi r\}^2 = 0$$

or

$$|r^H \Phi r + l^H \Phi l| = |2l^H \Psi r|$$

since $r^H \Phi r$ and $l^H \Phi l$ are always real.

To relate this to the conditioning the problem of finding $\hat{A} \in \mathcal{D}$ near A, it is useful delve more deeply into the geometry of $\mathcal{D}$. We will find that the matrices with ill-conditioned minima on $\mathcal{D}$ are actually sitting "focal points" of $\mathcal{D}$ analogous to focal points in geometric optics (explored in the next section).

It will be useful to first do some dimension counting in both complex and real versions of $\mathcal{D}$ so that we might better understand where the degrees of freedom are. (We shall always refer to one real parameter as one degree of freedom.)

The solution to the problem of finding a stationary point $\hat{A}$ to A can be expressed as decompositions of A as follows:

$$\begin{aligned} A &= suv^H + \hat{A}, \\ A &= suv^H + \lambda I + \tilde{U}\tilde{\Sigma}\tilde{V}^H, \\ A &= \lambda I + U\Sigma V^H \end{aligned}$$

where $u^H v = 0$, $\tilde{\Sigma} = \mathrm{diag}(\sigma_1, \ldots, \sigma_{n-1})$, $\Sigma = \mathrm{diag}(\sigma_1, \ldots, \sigma_n) = \begin{pmatrix} \tilde{\Sigma} & 0 \\ 0 & s \end{pmatrix}$, $U = (\tilde{U} \quad u)$, $V = (\tilde{V} \quad v)$, and U, V are unitary matrices. For the real case, this decomposes the n^2 degrees of freedom of A into $n(n-1)/2$ for U, $n(n-1)/2$ for V, n for Σ, and 1 for λ, which leaves exactly 1 degree of freedom extra to satisfy the constraint $u^H v = 0$. For the complex case, A have $2n^2$ degrees of freedom, U has n^2, V has n^2, Σ has n, and λ has 2, leaving $n+2$ surplus degrees of freedom. Since only the decomposition is only unique up to transformation of the form $U, V \to UZ, VZ$ (Z a diagonal unitary matrix), n of those surplus degrees of freedom are ineffective, leaving 2 to satisfy the constraints $\mathrm{Re}\{u^H v\} = 0, \mathrm{Im}\{u^H v\} = 0$.

For any element of $\mathcal{D}$, we can find a decomposition as above such that $s = 0$. In the real case, this eliminates one degree of freedom corresponding to $\mathcal{D}$'s real codimension of 1. In the complex case, this eliminates not only the degree of freedom in s, but also the extra unitary freedom of u and v corresponding to $\hat{A}$'s complex codimension of 1.

Now consider the normal points passing through a matrix $\hat{A} \in \mathcal{D}$. Using the above decomposition, we can parameterize this set as either a real or complex line: $A = \hat{A} + suv^H$, where s is real and the phases on u and v are arbitrary (alternatively, we could fix the phase of uv^H and let s be complex). This line has one degree of freedom in the reals and two in the complex. For most A we expect this decomposition to be locally unique. We may then define a "focal point" as some matrix $A = \hat{A} + suv^H$ as a matrix for which the decomposition is not unique, i.e. Jacobian of this decomposition is singular. More rigorously, we have:

Definition 10.1 *A matrix A is a "focal point" of the normal line $A = \hat{A} + suv^H$ when one can find nonzero $\delta\hat{A} \in T_{\mathcal{D}}\hat{A}$, δs, δu, and δv such that $\delta\hat{A} + \delta s uv^H + s\delta uv^H + su\delta v^H = 0$.*

We will now proceed to identify those points and show that these points are the set on which $\kappa = \infty$.

Varying the decomposition we have

$$\delta\lambda I + \delta U\Sigma V^H + U\delta\Sigma V^H + U\Sigma\delta V^H = 0$$

where $\delta u^H v + u^H\delta v = 0$. We will proceed by solving for an expression for $\delta u^H v + u^H\delta v$ in terms of U, S, V and then finding values for s and the phase of uv^H for which the variation of the constraint is satisfied (it is convenient for us to take s to be real and to adjust the phases of uv^H implicitly).

Multiplying by U^H and V on the right and left respectively, we have

$$\delta\lambda U^H V + (U^H\delta U)\Sigma + \Sigma(V^H\delta V)^H + \delta\Sigma = 0.$$

If $\sigma_n = 0$, then we may set $\delta\sigma_n$ to a nonzero value for a nontrivial solution of this equation. However, this is just the case when $A \in \mathcal{D}$.

We will assume that $\sigma_i > 0$ for all i. Transposing, multiplying by Σ, and some canceling terms gives:

$$\begin{aligned}
\Sigma^2(V^H\delta V) - (V^H\delta V)\Sigma^2 &= 2\delta\Sigma\Sigma + \delta\lambda\Sigma U^H V + \delta\bar{\lambda}V^H U\Sigma \\
\Sigma^2(U^H\delta U) - (U^H\delta U)\Sigma^2 &= 2\delta\Sigma\Sigma + \delta\bar{\lambda}\Sigma V^H U + \delta\lambda U^H V\Sigma
\end{aligned}$$

Note, the diagonals of the left hand sides are both zero. The diagonals of the right hand sides are both real.

Taking diagonals we have

$$\delta\Sigma = \mathrm{Re}\{\lambda \mathrm{diag}(U^H V)\}.$$

When $\delta\lambda = 0$, we have $\delta\Sigma = 0$. We would then have

$$\begin{aligned}
\Sigma^2(V^H\delta V) - (V^H\delta V)\Sigma^2 &= 0 \\
\Sigma^2(U^H\delta U) - (U^H\delta U)\Sigma^2 &= 0
\end{aligned}$$

which has nontrivial iff solutions when $\sigma_k = \sigma_{k+1}$ for some k (i.e. if the singular values cross).

Assuming that $\delta\lambda \neq 0$, however, we can always find $\delta\Sigma$ satisfying the diagonal equations.

We turn our attention to solving the off diagonal equations. Off the diagonal, we have:

$$\begin{aligned}
\Sigma^2(V^H\delta V) - (V^H\delta V)\Sigma^2 &= \delta\lambda\Sigma U^H V + \delta\bar{\lambda}V^H U\Sigma \\
\Sigma^2(U^H\delta U) - (U^H\delta U)\Sigma^2 &= \delta\bar{\lambda}\Sigma V^H U + \delta\lambda U^H V\Sigma
\end{aligned}$$

The operator of the left hand side is invertible. Since the right hand sides of both equations are hermitian, $V\delta V$ and $U\delta U$ will be anti-Hermitian as required.

It now remains to satisfy the varied constraint. $\delta u^H v + u^H\delta v = 0$. As before, we may rewrite this as $(\tilde{U}^H\delta u)^H r + l^H(\tilde{V}^H\delta v) = 0$.

Solving the off-diagonal equations for $\tilde{U}^H \delta u$ and $\tilde{V}^H \delta v$, we have

$$\tilde{V}^H \delta v = -\delta\lambda \Psi r - \delta\bar{\lambda}\Phi l$$
$$\tilde{U}^H \delta u = -\delta\bar{\lambda}\Psi l - \delta\lambda \Phi r.$$

Substituting into the constraint, gives

$$\delta\lambda(2l^H \Psi r) + \delta\bar{\lambda}(l^H \Phi l + r^H \Phi r) = 0.$$

For real λ this reduces to

$$2l^H \Psi r + l^H \Phi l + r^H \Phi r = 0,$$

i.e. $\sigma_{xx} = 0$. For complex λ, we can always select set the phase for of $\delta\lambda$ so that the two terms in the expression above have the same phase. We then have the condition on the parameter s as,

$$|2l^H \Psi r| = |l^H \Phi l + r^H \Phi r|.$$

Summarizing, we have found that the decomposition which reveals the nearby double eigenvalue structure of a generic matrix A is degenerate whenever $s = \sigma_{n-1}$ (i.e. nonsimple singular value) or s is a zero of $|2l^H \Psi r| - |l^H \Phi l + r^H \Phi r|$ (i.e. $\det \begin{pmatrix} \sigma_{xx} & \sigma_{xy} \\ \sigma_{yx} & \sigma_{yy} \end{pmatrix} = 0$). Geometrically, we may say that along the line passing through $\hat{A}$, there are certain s giving focal points where the normals from other points in $\mathcal{D}$ infinitesimally close to $\hat{A}$ intersect.

11 Principle Curvatures and Focal Points of $\mathcal{D}$

In this section we explore the differential geometry of $\mathcal{D}$ itself, as a co-dimension 1 manifold embedded in R^{n^2} or C^{n^2}. We will give an additional interpretation of the focal points of $\mathcal{D}$ in terms of principle curvatures.

The principle curvatures is typically studied in the setting of surfaces in R^3. In this setting, one uses the unit normal field $\hat{n}$ of a surface M to construct a map from M to S^2. Studying the local behavior of this map ultimately results in a quadratic form on the tangent space of M at a point. The reciprocal lengths of the principle axes of this quadratic form are called the principle curvatures, their sum called the mean curvature and their product the Gaussian curvature.

One way to generalize this approach which is suitable for any manifold of co-dimension 1 is to use the normal to construct a natural quadratic form on the tangent space which results from metrical considerations and then examining the eigenvalues of this form.

For practice, let's first look at an n-sphere of radius l embedded in R^{n+1}. Geodesics on S^n satisfy the equations

$$\begin{aligned}\frac{dp}{dt} &= v\\ \frac{dv}{dt} &= -\frac{v\cdot v}{l}\hat{n}\end{aligned}$$

where p is a point on the sphere, v satisfies $p^H v = 0$, and $\hat{n} = p/l$. One might recognize the derivative of v as just the centripetal acceleration in classical mechanics.

For a more general manifold of codimension 1, one has (at least locally) a unit normal $\hat{n}$, and geodesic equations of the form

$$\begin{aligned}\frac{dp}{dt} &= v\\ \frac{dv}{dt} &= -\Gamma(v,v)\end{aligned}$$

where $\Gamma(v,w)$ is the Levi-Civita connection. However, for manifolds of codimension 1, the Levi-Civita connection must be of the form $\Gamma(v,w) = a(v,w)\hat{n}$ where the symmetric bilinear function $a(v,w)$ defined on the tangent space at p is determined by $\frac{dv}{dt}\cdot\hat{n} + v\cdot\frac{d\hat{n}}{dt} = 0$.

If we call $a(v,v)$ the centripetal acceleration along a geodesic with velocity v through p, then we may define the principle curvatures at p as the extrema of the function $|a(v,v)|$ over the unit ball in the tangent space at p. Like the sphere, center of curvature is located at $p - \frac{1}{a(v,v)}\hat{n}$.

With this in mind, we now consider the manifold $\mathcal{D}$. For the remainder of this section any matrix A is a point in $\mathcal{D}$ with decomposition

$$A = \tilde{U}\tilde{\Sigma}\tilde{V}^H + \lambda I$$

and normal $N = uv^H$ $(u^H v = 0)$. Additionally $r = \tilde{U}^H v$ and $l = \tilde{V}^H u$. Any element H of the tangent space of A must satisfy $u^H H v = 0$ i.e. must be perpendicular to the normal.

Thus any element of the tangent space of $\mathcal{D}$ at A has the form

$$H = \tilde{U}E\tilde{V}^H + \tilde{U}yv^H + ux^H\tilde{V}^H$$

where E, x, y are arbitrary.

If δA is a first order change in the decomposition of A then δA can be written as

$$\delta A = \delta\tilde{U}\tilde{\Sigma}\tilde{V}^H + \tilde{U}\delta\tilde{\Sigma}\tilde{V}^H + \delta\tilde{U}\tilde{\Sigma}\delta\tilde{V}^H + \mu I$$

where

$$\begin{aligned} \delta\tilde{U} &= \tilde{U}J_1 + uj_1^H \\ \delta\tilde{V} &= \tilde{V}J_2 + vj_2^H \\ \delta\tilde{\Sigma} &= D \\ l^H\delta\tilde{V}^Hv + u^H\delta\tilde{U}r &= l^Hj_2 + j_1^Hr = 0 \end{aligned}$$

for J_i antisymmetric, j_i $n-1$ vectors, and D diagonal. The last equation results from varying the equation $u^Hv = 0$. Collecting terms

$$\delta A = \delta\tilde{U}(J_1\tilde{\Sigma} - \tilde{\Sigma}J_2 + D)\tilde{V}^H + uj_1^H\tilde{\Sigma}\tilde{V}^H + \tilde{U}\tilde{\Sigma}j_2v^H + \mu I.$$

Equating $\delta A = H$ and using orthogonality we have the following relations

$$\begin{array}{rcl} \tilde{U}^HH\tilde{V} &=& (J_1\tilde{\Sigma} - \tilde{\Sigma}J_2 + D) + \mu\tilde{U}^H\tilde{V} = E \\ u^HH\tilde{V} &=& j_1^H\tilde{\Sigma} + \mu l^H = x^H \\ \tilde{U}^HHv &=& \tilde{\Sigma}j_2 + \mu r = y \end{array}$$

The varied constraint equation becomes

$$l^Hj_2 + j_1^Hr = l^H\tilde{\Sigma}^{-1}y + x^H\tilde{\Sigma}^{-1}r - 2\mu l^H\tilde{\Sigma}^{-1}r = 0$$

which can be used to solve for μ, which then fixes J_1, J_2, and D (though we will not be concerned with expressions for them), as well as

$$\begin{aligned} j_1^H &= (x^H - \mu l^H)\tilde{\Sigma}^{-1} \\ j_2 &= \tilde{\Sigma}^{-1}(y - \mu r) \end{aligned}$$

which will will have occasion to use.

To determine $a(H,H)$ we may use $a(H,H) = \delta u^H H v + u^H H \delta v$, which gives

$$\begin{aligned}
a(H,H) &= -u^H \delta\tilde{U}\tilde{U}^H H v - u^H H \tilde{V}\delta\tilde{V}^H v \\
a(H,H) &= -j_1^H y - x^H j_2 \\
a(H,H) &= -(x^H - \mu l^H)\tilde{\Sigma}^{-1} y - x^H \tilde{\Sigma}^{-1}(y - \mu r) \\
a(H,H) &= -2x^H \tilde{\Sigma}^{-1} y + \frac{(l^H \tilde{\Sigma}^{-1} y + x^H \tilde{\Sigma}^{-1} r)^2}{2 l^H \tilde{\Sigma}^{-1} r}
\end{aligned}$$

We see that the quadratic form $a(H,H)$ can be written as $z^t W z$ where

$$z = \begin{pmatrix} \bar{x} \\ y \end{pmatrix}$$

and

$$W = \begin{pmatrix} 0 & -\tilde{\Sigma}^{-1} \\ -\tilde{\Sigma}^{-1} & 0 \end{pmatrix} + \frac{1}{2\bar{l}^t \tilde{\Sigma}^{-1} r} \begin{pmatrix} \tilde{\Sigma}^{-1} r \\ \tilde{\Sigma}^{-1} \bar{l} \end{pmatrix} \begin{pmatrix} r^t \tilde{\Sigma}^{-1} & \bar{l}^t \tilde{\Sigma}^{-1} \end{pmatrix}.$$

By a simple orthogonal similarity transformation we have

$$QWQ^H = \begin{pmatrix} -\tilde{\Sigma}^{-1} & 0 \\ 0 & \tilde{\Sigma}^{-1} \end{pmatrix} + \frac{1}{4\bar{l}^t \tilde{\Sigma}^{-1} r} \begin{pmatrix} \tilde{\Sigma}^{-1}(r + \bar{l}) \\ \tilde{\Sigma}^{-1}(r - \bar{l}) \end{pmatrix} \begin{pmatrix} (r + \bar{l})^t \tilde{\Sigma}^{-1} & (r - \bar{l})^t \tilde{\Sigma}^{-1} \end{pmatrix},$$

a real diagonal matrix plus a rank-1 matrix.

We now seek the extrema of $z^t W z$ where $|H|^2 = |x|^2 + |y|^2 = 1$. Note that $W^t = W$. Varying $|z^t W z|$ we see that the condition for z to be an extrema of $|z^t W z|$ is

$$Wz = \zeta \bar{z},$$

a somewhat unusual eigenvalue equation over C.

Over R this reduces to the problem of finding the eigenvalues of a rank-1 update, $D + aa^t$. The eigenvalues, ζ, of this matrix then satisfy

$$a^t(\zeta I - D)^{-1} a = 1,$$

which, in our terms, is

$$\begin{pmatrix} (r+l)^t \tilde{\Sigma}^{-1} & (r-l)^t \tilde{\Sigma}^{-1} \end{pmatrix} \begin{pmatrix} (\zeta I + \tilde{\Sigma}^{-1})^{-1} & 0 \\ 0 & (\zeta I - \tilde{\Sigma}^{-1})^{-1} \end{pmatrix} \begin{pmatrix} \tilde{\Sigma}^{-1}(r+l) \\ \tilde{\Sigma}^{-1}(r-l) \end{pmatrix} = 4 l^t \tilde{\Sigma}^{-1} r$$

multiplying this out and simplifying gives,

$$r^t \frac{-1}{\zeta}\left(\frac{1}{\zeta^2} I - \tilde{\Sigma}^2\right)^{-1} r + l^t \frac{-1}{\zeta}\left(\frac{1}{\zeta^2} - \tilde{\Sigma}^2\right)^{-1} l + 2 l^t \tilde{\Sigma}\left(\frac{1}{\zeta^2} - \tilde{\Sigma}^2\right)^{-1} r = 0.$$

$$r^t \Phi\left(\frac{-1}{\zeta}\right) r + l^t \Phi\left(\frac{-1}{\zeta}\right) l + 2 l^t \Psi\left(\frac{-1}{\zeta}\right) r = 0,$$

where $\Phi(s) = s(s^2 I - \tilde{\Sigma}^2)^{-1}$ and $\Psi(s) = \tilde{\Sigma}(s^2 I - \tilde{\Sigma}^2)^{-1}$. The center of curvature is then $A + suv^H$ where $s = \frac{-1}{\zeta}$, we may now interpret the vanishing of the second real-axis derivative at a stationary point of the spectral portrait as equivalent to being on a focal point of the real double eigenvalue manifold.

Over C, we solve $QWQ^H(Qz) = \zeta Q\bar{z}$ by first observing that given an eigen-pair z,ζ , the eigen-pair $ze^{i\frac{1}{2}\angle\zeta}$, $\zeta e^{-i\angle\zeta}$ is also a solution. Thus, while acknowledging that λ can actually have any phase. we may restrict ζ to the reals for purposes of computation. In which case we solve,

$$\begin{pmatrix} QWQ^H & -\zeta I \\ -\zeta I & Q\bar{W}Q^H \end{pmatrix} \begin{pmatrix} z \\ \bar{z} \end{pmatrix}.$$

Following a derivation of the usual rank-1 update formula, we rewrite this as

$$\begin{pmatrix} D & -\zeta I \\ -\zeta I & D \end{pmatrix} \begin{pmatrix} z \\ \bar{z} \end{pmatrix} + \begin{pmatrix} aa^t & 0 \\ 0 & \bar{a}\bar{a}^t \end{pmatrix} \begin{pmatrix} z \\ \bar{z} \end{pmatrix} = 0$$

where

$$D = \begin{pmatrix} -\tilde{\Sigma}^{-1} & 0 \\ 0 & \tilde{\Sigma}^{-1} \end{pmatrix},$$

$$a = \frac{1}{2\sqrt{\bar{l}^t\tilde{\Sigma}^{-1}r}} \begin{pmatrix} \tilde{\Sigma}^{-1}(r+\bar{l}) \\ \tilde{\Sigma}^{-1}(r-\bar{l}) \end{pmatrix}$$

Letting $k = a^t z$ we have

$$\begin{pmatrix} D & -\zeta I \\ -\zeta I & D \end{pmatrix} \begin{pmatrix} z \\ \bar{z} \end{pmatrix} + \begin{pmatrix} ak \\ \bar{a}\bar{k} \end{pmatrix} = 0$$

$$\begin{pmatrix} z \\ \bar{z} \end{pmatrix} + \begin{pmatrix} D(D^2-\zeta^2 I)^{-1}a & \zeta(D^2-\zeta^2 I)^{-1}\bar{a} \\ \zeta(D^2-\zeta^2 I)^{-1}a & D(D^2-\zeta^2 I)^{-1}\bar{a} \end{pmatrix} \begin{pmatrix} k \\ \bar{k} \end{pmatrix}$$

$$\begin{pmatrix} k \\ \bar{k} \end{pmatrix} + \begin{pmatrix} a^t D(D^2-\zeta^2 I)^{-1}a & a^t\zeta(D^2-\zeta^2 I)^{-1}\bar{a} \\ \bar{a}^t\zeta(D^2-\zeta^2 I)^{-1}a & \bar{a}^t D(D^2-\zeta^2 I)^{-1}\bar{a} \end{pmatrix} \begin{pmatrix} k \\ \bar{k} \end{pmatrix}.$$

Note that the off diagonals are real, while the diagonals are each others complex conjugates, and that $\begin{pmatrix} k \\ \bar{k} \end{pmatrix}$ must be a fixed point of this matrix. A necessary and sufficient condition for such a k to exist is

$$\det \begin{pmatrix} a^t D(D^2-\zeta^2 I)^{-1}a + 1 & a^t\zeta(D^2-\zeta^2 I)^{-1}\bar{a} \\ \bar{a}^t\zeta(D^2-\zeta^2 I)^{-1}a & \bar{a}^t D(D^2-\zeta^2 I)^{-1}\bar{a} + 1 \end{pmatrix} = 0,$$

or equivalently,

$$|a^t D(D^2-\zeta^2 I)^{-1}a + 1| = |\bar{a}^t\zeta(D^2-\zeta^2 I)^{-1}a|.$$

Substituting and simplifying we have

$$|\bar{r}^t\Phi(\frac{-1}{\zeta})r + \bar{l}^t\Phi(\frac{-1}{\zeta})l| = |2\bar{l}^t\Psi(\frac{-1}{\zeta})r|.$$

Thus we see that an additional way to interpret these focal points is to imagine a circle of radius $\frac{1}{\zeta}$, centered at the focal point and passing through $\mathcal{D}$ along a principle direction. This circle then passes through $\mathcal{D}$ with second order contact [].