NUMERICAL METHODS FOR SOLVING INVERSE EIGENVALUE PROBLEMS FOR NONNEGATIVE MATRICES*

ROBERT ORSI[†]

Abstract. Presented here are two related numerical methods, one for the inverse eigenvalue problem for nonnegative or stochastic matrices and another for the inverse eigenvalue problem for symmetric nonnegative matrices. The methods are iterative in nature and utilize alternating projection ideas. For the algorithm for the symmetric problem, the main computational component of each iteration is an eigenvalue-eigenvector decomposition, while for the other algorithm, it is a Schur matrix decomposition. Convergence properties of the algorithms are investigated and numerical results are also presented. While the paper deals with two specific types of inverse eigenvalue problems, the ideas presented here should be applicable to many other inverse eigenvalue problems, including those involving nonsymmetric matrices.

Key words. inverse eigenvalue problem, nonnegative matrices, stochastic matrices, alternating projections, Schur's decomposition

AMS subject classifications. 15A51, 65F18

DOI. 10.1137/050634529

1. Introduction. A real $n \times n$ matrix is said to be nonnegative if each of its entries is nonnegative.

The *nonnegative inverse eigenvalue problem* (NIEP) is the following: given a list of n complex numbers $\lambda = \{\lambda_1, \dots, \lambda_n\}$, find a nonnegative $n \times n$ matrix with eigenvalues λ (if such a matrix exists).

A related problem is the *symmetric nonnegative inverse eigenvalue problem* (SNIEP): given a list of n real numbers $\lambda = \{\lambda_1, \dots, \lambda_n\}$, find a symmetric nonnegative $n \times n$ matrix with eigenvalues λ (if such a matrix exists)[1].

Finding necessary and sufficient conditions for a list λ to be realizable as the eigenvalues of a nonnegative matrix has been a challenging area of research for over fifty years, and this problem is still unsolved [12]. As noted in [6, section 6], while various necessary or sufficient conditions exist, the necessary conditions are usually too general while the sufficient conditions are too specific. Under a few special sufficient conditions, a nonnegative matrix with the desired spectrum can be constructed; however, in general, proofs of sufficient conditions are nonconstructive. Two sufficient conditions that are constructive and not restricted to small n are, respectively, given in [20], for the SNIEP, and [21], for the NIEP with real λ. (See also [19] for an extension of the results of the latter paper.) A good overview of known results relating to necessary or sufficient conditions can be found in the recent survey paper [12] and general background material on nonnegative matrices, including inverse eigenvalue problems and applications, can be found in the texts [2] and [18]. We also mention the recent paper [9], which can be used to help determine whether a given list λ may be realizable as the eigenvalues of a nonnegative matrix.

†Research School of Information Sciences and Engineering, Australian National University, Canberra ACT 0200, Australia (robert.orsi@anu.edu.au).

[1]The NIEP and SNIEP are different problems even if λ is restricted to contain only real entries; there exist lists of n real numbers λ for which the NIEP is solvable but the SNIEP is not [16].

In this paper we are interested in generally applicable numerical methods for solving NIEPs and SNIEPs. To the best of our knowledge, the only algorithms that have appeared up to now in the literature consist of [5] for the SNIEP and [8] for the NIEP. In the case of [5], the following constrained optimization problem is considered:

$$(1.1) \qquad \min_{Q^T Q=I, \ R=R^T} \frac{1}{2}\|Q^T \Lambda Q - R \circ R\|^2.$$

Here Λ is a constant diagonal matrix with the desired spectrum and $\circ$ stands for the Hadamard product, i.e., componentwise product. Note that the symmetric matrices with the desired spectrum are exactly the elements of $\{Q^T \Lambda Q \mid Q \in \mathbb{R}^{n \times n} \text{ orthogonal}\}$ and that the symmetric nonnegative matrices are exactly the elements of $\{R \circ R \mid R \in \mathbb{R}^{n \times n} \text{ symmetric}\}$. In [5], a gradient flow based on (1.1) is constructed. A solution to the SNIEP is found if the gradient flow converges to a Q and an R that zero the objective function. The approach taken in [8] for the NIEP is similar but is complicated by the fact that the set of all matrices, both symmetric and nonsymmetric, with a particular desired spectrum is not nicely parameterizable. In particular, these matrices can no longer be parameterized by the orthogonal matrices.

In this paper we present a numerical algorithm for the NIEP and another for the SNIEP. In both cases, the problems are posed as problems of finding a point in the intersection of two particular sets. Unlike the approaches in [5] and [8] which are based on gradient flows, our algorithms are iterative in nature. For the SNIEP, the solution methodology is based on a alternating projection scheme between the two sets in question. The solution methodology for the NIEP is also based on an alternating projection-like scheme but is more involved, as we will shortly explain.

While alternating projections can often be a very effective means of finding a point in the intersection of two or more convex sets, for both the SNIEP and NIEP formulations, one set is nonconvex. Nonconvexity of one of the sets means that alternating projections may not converge to a solution. This is in contrast to the case where all sets are convex and convergence to a solution is guaranteed.

In addition to problem formulations, the development of the corresponding algorithms, and their convergence analysis, another contribution of the paper is as follows. As mentioned above, for each problem, one set in the problem formulation is nonconvex. For the NIEP, this set is particularly complicated; it consists of all matrices with the desired spectrum. At least some of the members of this set will be nonsymmetric matrices and it is this that causes complications. In particular, though the set is closed and hence projections are well defined theoretically, how to calculate projections onto such sets is an unsolved difficult problem. We formulate an alternate method for mapping onto this set. Though the resulting points are not necessarily projected points, they are members of the set and share a number of other desirable properties. As will be shown, this alternate "projection" is very effective in our context. Furthermore, we believe that it may also be quite effective for other inverse eigenvalue problems involving nonsymmetric matrices[2]. For more on other inverse eigenvalue problems, see the survey papers [4] and [6], and the recent text [7].

Before concluding this introductory section we would like to point out how the NIEP is related to another problem involving stochastic matrices. A $n \times n$ matrix is said to be stochastic if it is nonnegative and the sum of the entries in each row equals one. Another variation of the NIEP is the *stochastic inverse eigenvalue problem*

---

[2]Preliminary indications of this are given in [24] and [23], where this idea is applied to inverse eigenvalue type problems arising in control theory.

(StIEP): given a list of n complex numbers $\lambda = \{\lambda_1, \ldots, \lambda_n\}$, find a stochastic $n \times n$ matrix with eigenvalues λ (if such a matrix exists). It turns out that the NIEP and the StIEP are almost exactly the same problem, as we now show. (See also [8].)

The vector of all 1's is always an eigenvector for a stochastic matrix, implying each stochastic matrix must have 1 as an eigenvalue. Also, the maximum row sum matrix norm of a stochastic matrix equals 1 and hence the spectral radius cannot be greater than 1, and as a result, must actually equal 1. Suppose λ satisfies the above mentioned necessary conditions to be the spectrum of a stochastic matrix and that a nonnegative matrix A with this spectrum can be found. Then if an eigenvector x of A corresponding to the eigenvalue 1 can be chosen to have positive entries (by the Perron–Frobenius theorem this is certainly possible if A is irreducible), then, if we define $D = \mathrm{diag}(x)$, it is straightforward to verify that

$$D^{-1}AD$$

is a stochastic matrix with the desired spectrum. (In fact it can be shown that if λ satisfies the above mentioned necessary conditions, then it is the spectrum of a stochastic matrix if and only if it is the spectrum of a nonnegative matrix [22, Lemma 5.3.2].)

The rest of the paper is structured as follows. The last part of this section contains some notation. Projections play a key part in the algorithms and section 2 contains general properties of projections that are used throughout the paper. The SNIEP algorithm is presented first, in section 3, and then insights from this algorithm are used to address the more difficult NIEP in section 4. Section 5 contains convergence results. This includes a detailed analysis of fixed points of the SNIEP algorithm for the $n = 2$ case. This is the easiest case though we believe the analysis presented is still quite interesting and also gives insight into higher-dimensional problems. Numerical results for both algorithms are presented in section 6, and an appendix contains some supplementary projection results.

Notation. $\mathbb{R}$ is the set of real numbers. $\mathbb{C}$ is the set of complex numbers. $\mathcal{S}^n$ is the set of real symmetric $n \times n$ matrices. A^T denotes the transpose of a matrix A. A^* denotes the complex conjugate transpose of a matrix A. $\mathrm{tr}(A)$ denotes the sum of the diagonal elements of a square matrix A. For two $n \times n$ symmetric matrices A and B, $[A, B]$ denotes $AB - BA$. $\mathrm{diag}(v)$ for $v \in \mathbb{C}^n$ denotes the $n \times n$ diagonal matrix whose ith diagonal term is v_i. $\mathrm{Re}(z)$ denotes the real part of $z \in \mathbb{C}$.

2. Projections. Projections play a key part in the algorithms. This section contains general properties of projections that will be used throughout the paper.

Let x be an element in a Hilbert space H and let C be a closed (possibly nonconvex) subset of H. Any $c_0 \in C$ such that $\|x - c_0\| \leq \|x - c\|$ for all $c \in C$ will be called a *projection* of x onto C. In the cases of interest here, namely where H is a finite dimensional Hilbert space, there is always at least one such point for each x. If C is convex as well as closed, then each x has exactly one such minimum distance point [17]. Where convenient, we will use $y = P_C(x)$ to denote that y is a projection of x onto C. We emphasize that $y = P_C(x)$ only says y is *a* projection of x onto C and does *not* make any statement regarding uniqueness.

All problems of interest in this paper are feasibility problems of the following abstract form.

PROBLEM 2.1. *Given closed sets $C_1, \ldots, C_N$ in a finite dimensional Hilbert space*

H, find a point in the intersection

$$\bigcap_{i=1}^{N} C_i$$

(assuming the intersection is nonempty).

(In fact, we will solely be interested in the case $N = 2$.)

If all the C_i's in Problem 2.1 are convex, a classical method of solving Problem 2.1 is to alternatively project onto the C_i's. This method is often referred to as the method of alternating projections (MAP). If the C_i's have a nonempty intersection, the successive projections are guaranteed to asymptotically converge to an intersection point [3].

THEOREM 2.2 (MAP). *Let $C_1, \ldots, C_N$ be closed convex sets in a finite dimensional Hilbert space H. Suppose $\bigcap_{i=1}^{N} C_i$ is nonempty. Then starting from an arbitrary initial value x_0, the sequence*

$$x_{i+1} = P_{C_{\phi(i)}}(x_i), \ \text{where } \phi(i) = (i \bmod N) + 1,$$

converges to an element in $\bigcap_{i=1}^{N} C_i$.

We remark that the usefulness of MAP for finding a point in the intersection of a number of sets is dependent on being able to compute projections onto each of the C_i's.

While MAP is not guaranteed to converge to a solution if one or more of the C_i's is nonconvex, for alternating projections between two sets, the following distance reduction property always holds.

THEOREM 2.3. *Let C_1 and C_2 be closed (nonempty) sets in a finite dimensional Hilbert space H. For any initial value $y_0 \in C_2$, if*

$$x_1 = P_{C_1}(y_0), \quad y_1 = P_{C_2}(x_1), \quad x_2 = P_{C_1}(y_1),$$

then

$$\|x_2 - y_1\| \le \|x_1 - y_1\| \le \|x_1 - y_0\|.$$

Proof. The second inequality holds as y_1 is a projection of x_1 onto C_2 and hence its distance to x_1 is less than or equal to the distance of x_1 to any other point in C_2 such as y_0. The first inequality holds by similar reasoning.   □

COROLLARY 2.4. *If for $i = 0, 1, \ldots$,*

$$x_{i+1} = P_{C_1}(y_i), \quad y_{i+1} = P_{C_2}(x_{i+1}),$$

that is, the x_i's and y_i's are successive projections between two closed sets, then $\|x_i - y_i\|$ is a nonincreasing function of i.

Suppose one is interested in solving Problem 2.1 in the case of two sets, C_1 and C_2, when one or both sets are nonconvex. If projections onto these sets are computable, a solution method is to alternately project onto C_1 and C_2. Corollary 2.4 ensures that the distance $\|x_i - y_i\|$ is nonincreasing with i. While this is promising, there is, however, no guarantee that this distance goes to zero and hence that a solution to the problem will be found.

Most of the literature on alternating projection methods deals with the case of convex subsets of a (possibly infinite dimensional) Hilbert space; a survey of these results is contained in [1]. The text [11] is also recommended. There is much less available for the case of one or more nonconvex sets; see in particular [10].

3. The symmetric problem. Our algorithm for solving the SNIEP consists of alternately projecting onto two particular sets. The details are given in this section.

Given a list of real eigenvalues $\lambda = \{\lambda_1, \dots, \lambda_n\}$, renumbering if necessary, suppose $\lambda_1 \geq \cdots \geq \lambda_n$. Let

$$(3.1) \qquad \Lambda = \operatorname{diag}(\lambda_1, \dots, \lambda_n),$$

and let $\mathcal{M}$ denote the set of all real symmetric matrices with eigenvalues λ,

$$(3.2) \qquad \mathcal{M} = \{A \in \mathcal{S}^n \mid A = V\Lambda V^T \text{ for some orthogonal } V\}.$$

Let $\mathcal{N}$ denote the set of symmetric nonnegative matrices,

$$(3.3) \qquad \mathcal{N} = \{A \in \mathcal{S}^n \mid A_{ij} \geq 0 \text{ for all } i,j\}.$$

The SNIEP can now be stated as the following particular case of Problem 2.1:

$$(3.4) \qquad \text{Find } X \in \mathcal{M} \cap \mathcal{N}.$$

Our solution approach is to alternatively project between $\mathcal{M}$ and $\mathcal{N}$, and we next show that it is indeed possible to calculate projections onto these sets. First, in order for the term "projection" to make sense, we need to define an appropriate Hilbert space and associated norm. From now on, $\mathcal{S}^n$ will be viewed as a Hilbert space with inner product

$$(3.5) \qquad \langle A, B \rangle = \operatorname{tr}(AB) = \sum_{i,j} A_{ij} B_{ij}.$$

The associated norm is the Frobenius norm $\|A\| = \langle A, A\rangle^{\frac{1}{2}}$.

The projection of $A \in \mathcal{S}^n$ onto $\mathcal{M}$ is given by Theorem 3.2 below. More precisely, it gives *a* projection of A onto $\mathcal{M}$. The reason for this is that the set $\mathcal{M}$ is nonconvex[3] and hence projections onto this set are not guaranteed to be unique. We will need the following classical result [13, section 10.2].

LEMMA 3.1. *Suppose $x, y \in \mathbb{R}^n$ and $x_1 \geq \cdots \geq x_n$, $y_1 \geq \cdots \geq y_n$. If σ is a permutation of $\{1, \dots, n\}$, then*

$$\sum_i x_i y_i \geq \sum_i x_i y_{\sigma(i)}.$$

THEOREM 3.2. *Given $A \in \mathcal{S}^n$, let $A = V\operatorname{diag}(\mu_1, \dots, \mu_n)V^T$ with V a real orthogonal matrix and $\mu_1 \geq \cdots \geq \mu_n$. If Λ is given by (3.1), then $V\Lambda V^T$ is a best approximant in $\mathcal{M}$ to A in the Frobenius norm.*

Proof. For all $X \in \mathcal{M}$, $\operatorname{tr}(X^2) = \operatorname{tr}(\Lambda^2)$. As a result, finding $X \in \mathcal{M}$ that minimizes $\|X - A\|^2$ is the same as finding $X \in \mathcal{M}$ that maximizes $\operatorname{tr}(XA)$. Consider the function

$$f : \mathcal{M} \to \mathbb{R}, \ X \mapsto \operatorname{tr}(XA).$$

---

[3]$\mathcal{M}$ is nonconvex if its defining λ contains a pair of nonequal eigenvalues. For example, if $n = 2$, consider

$$A = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} \lambda_2 & 0 \\ 0 & \lambda_1 \end{bmatrix}.$$

If $\lambda_1 \neq \lambda_2$, then the convex combination $(A+B)/2$ does not have the same spectrum as A and B.

$\mathcal{M}$ is a smooth manifold and its tangent space at a point X is $\{[X, \Omega] \mid \Omega = -\Omega^T \in \mathbb{R}^{n \times n}\}$; see, for example, [14, Chapter 2]. The derivative of f at a point X in the tangent direction $[X, \Omega]$ is

$$Df(X)([X, \Omega]) = \text{tr}([X, \Omega]A) = \text{tr}((AX - XA)\Omega).$$

If X maximizes f, then this derivative must be zero in all tangent directions, or equivalently, $AX - XA$ must be symmetric. This in turn is equivalent to X and A commuting. X and A commute if and only if they are simultaneously diagonalizable; see [15, Theorem 2.5.15]. Hence if X maximizes f, then there must exist an orthogonal matrix U and a diagonal matrix Λ_σ with the same spectrum as Λ such that $A = U \text{diag}(\mu_1, \dots, \mu_n)U^T$ and $X = U\Lambda_\sigma U^T$. This combined with Lemma 3.1 implies f has maximum value $\text{tr}(\Lambda \text{diag}(\mu_1, \dots, \mu_n))$ and implies the result. $\square$

Projection onto $\mathcal{N}$ is straightforward and is given by Theorem 3.3 below.

THEOREM 3.3. *Given $A \in \mathcal{S}^n$, define $A_+ \in \mathcal{S}^n$ by*

(3.6) $$(A_+)_{ij} = \max\{A_{ij}, 0\} \text{ for all } 1 \le i, j \le n.$$

A_+ is the best approximant in $\mathcal{N}$ to A in the Frobenius norm.

Proof. The projection of $x \in \mathbb{R}$ onto the nonnegative real numbers equals $\max\{x, 0\}$. The general result follows by noting that if $B \in \mathcal{S}^n$, and in particular if $B \in \mathcal{N}$, then

$$\|A - B\| = \left(\sum_{i,j} |A_{ij} - B_{ij}|^2 \right)^{\frac{1}{2}}$$

and hence that the problem reduces to n^2 decoupled scalar problems. $\square$

Our proposed algorithm for solving the SNIEP is the following.

SNIEP algorithm:

Problem data. List of desired real eigenvalues $\lambda = \{\lambda_1, \dots, \lambda_n\}$, $\lambda_1 \ge \dots \ge \lambda_n$.

Initialization. Choose a randomly generated symmetric nonnegative matrix $Y \in \mathbb{R}^{n \times n}$.

repeat

    1. Calculate an eigenvalue-eigenvector decomposition of Y:
       $Y = V \text{diag}(\mu_1, \dots, \mu_n)V^T$, $\mu_1 \ge \dots \ge \mu_n$.
    2. $X := V \text{diag}(\lambda_1, \dots, \lambda_n)V^T$.
    3. $X := (X + X^T)/2$.
    4. $Y := X_+$.

  until $\|X - Y\| < \epsilon$.

In the above algorithm, X_+ is given by (3.6).

Note that at each iteration of the algorithm, X has the desired spectrum λ and Y is nonnegative. If ϵ is small, say $\epsilon = 10^{-14}$, termination of the loop ensures X equals Y (approximately) and hence that Y solves the SNIEP.

Due to small numerical inaccuracy, X from Step 2 of the algorithm may not be perfectly symmetric. Step 3 makes it so.

Of course, while Corollary 2.4 ensures $\|X - Y\|$ is nonincreasing from one iteration to the next, the set $\mathcal{M}$ is nonconvex and hence there is no guarantee that the algorithm will terminate. A detailed analysis of convergence is postponed to section 5.

4. The general problem. Throughout this section, $\mathbb{C}^{n \times n}$ will be viewed as a Hilbert space with inner product

$$\langle A, B \rangle = \operatorname{tr}(AB^*) = \sum_{i,j} A_{ij} \overline{B}_{ij}.$$

The associated norm is the Frobenius norm $\|A\| = \langle A, A \rangle^{\frac{1}{2}}$.

Recall Schur's result that any matrix $A \in \mathbb{C}^{n \times n}$ is unitarily equivalent to an upper triangular matrix.

THEOREM 4.1. *Given $A \in \mathbb{C}^{n \times n}$ with eigenvalues $\mu_1, \dots, \mu_n$ in any prescribed order, there is a unitary matrix $U \in \mathbb{C}^{n \times n}$ and an upper triangular matrix $T \in \mathbb{C}^{n \times n}$ such that*

$$(4.1) \qquad\qquad\qquad\qquad A = UTU^*$$

and $T_{ii} = \mu_i$, $i = 1, \dots, n$.

Proof. See, for example, [15, Theorem 2.3.1].   □

We now redefine some terms from the prior section.

Let $\lambda = \{\lambda_1, \dots, \lambda_n\}$ be a given list of complex eigenvalues. Define

$$(4.2) \qquad \mathcal{T} = \{T \in \mathbb{C}^{n \times n} \,|\, T \text{ is upper triangular with spectrum } \lambda\}.$$

Theorem 4.1 implies that the set of all complex matrices with spectrum λ is given by the following set:

$$(4.3) \qquad \mathcal{M} = \{A \in \mathbb{C}^{n \times n} \,|\, A = UTU^* \text{ for some unitary } U \text{ and some } T \in \mathcal{T}\}.$$

Let $\mathcal{N}$ denote the set of (not necessarily symmetric) nonnegative matrices,

$$(4.4) \qquad\qquad \mathcal{N} = \{A \in \mathbb{R}^{n \times n} \,|\, A_{ij} \geq 0 \text{ for all } i, j\}.$$

Having redefined $\mathcal{M}$ and $\mathcal{N}$, the NIEP can now be stated as the following particular case of Problem 2.1:

$$(4.5) \qquad\qquad\qquad\qquad \text{Find } X \in \mathcal{M} \cap \mathcal{N}.$$

A difficulty now occurs. We would like to use alternating projections to solve the NIEP. However, to the best of our knowledge, the way to calculate projections onto $\mathcal{M}$ is an unsolved problem. Suppose instead we could find a mapping that was in some sense a reasonable substitute for a projection map for $\mathcal{M}$. Using this substitute mapping and the projection map for $\mathcal{N}$ in an alternating projection-like scheme may still produce a viable algorithm. Indeed, we now propose the following function $P_{\mathcal{M}}$ as a substitute for a true projection map onto $\mathcal{M}$. (The notation $P_{\mathcal{M}}$ is used as it is suggestive; however, recall that we have already used $y = P_C(x)$ to denote that y is a projection of x onto a set C. The two different uses of the notation should be clear from their context and should not cause confusion.)

DEFINITION 4.2. *Suppose $U \in \mathbb{C}^{n \times n}$ is unitary and $T \in \mathbb{C}^{n \times n}$ is upper triangular. Let $\{\hat{\lambda}_1, \dots, \hat{\lambda}_n\}$ be a permutation of the list of eigenvalues λ such that, among all possible permutations, it minimizes*

$$(4.6) \qquad\qquad\qquad\qquad \sum_{i=1}^{n} |\hat{\lambda}_i - T_{ii}|^2.$$

Define

(4.7) $$P_{\mathcal{M}}(U,T) = U\hat{T}U^*,$$

where $\hat{T} \in \mathcal{T}$ is given by

$$\hat{T}_{ij} = \begin{cases} \hat{\lambda}_i & \text{if } i = j, \\ T_{ij} & \text{otherwise.} \end{cases}$$

Note that $P_{\mathcal{M}}$ maps into the set $\mathcal{M}$.

A given matrix $A \in \mathbb{C}^{n \times n}$ may have a nonunique Schur decomposition and $A = U_1 T_1 U_1^* = U_2 T_2 U_2^*$ does *not* imply $P_{\mathcal{M}}(U_1, T_1) = P_{\mathcal{M}}(U_2, T_2)$. For example, if

$$T_1 = \begin{bmatrix} 1 & 1 & 4 \\ 0 & 2 & 2 \\ 0 & 0 & 3 \end{bmatrix}, \quad T_2 = \begin{bmatrix} 2 & -1 & 3\sqrt{2} \\ 0 & 1 & \sqrt{2} \\ 0 & 0 & 3 \end{bmatrix}, \quad \text{and } U = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 & 0 \\ 1 & -1 & 0 \\ 0 & 0 & \sqrt{2} \end{bmatrix},$$

then U is unitary and $UT_1U^* = T_2$, [15]. If $\lambda = \{0,0,0\}$, $P_{\mathcal{M}}(U, T_1) \neq P_{\mathcal{M}}(I, T_2)$.

It turns out that this nonuniqueness is not particularly important. The following result shows that for different Schur decompositions of the same matrix, $P_{\mathcal{M}}$ gives points in $\mathcal{M}$ of equal distance from the original matrix.

THEOREM 4.3. *Suppose $A = U_1 T_1 U_1^* = U_2 T_2 U_2^*$, where $U_1, U_2 \in \mathbb{C}^{n \times n}$ are unitary and $T_1, T_2 \in \mathbb{C}^{n \times n}$ are upper triangular. Then*

$$\|P_{\mathcal{M}}(U_1, T_1) - A\| = \|P_{\mathcal{M}}(U_2, T_2) - A\|.$$

Proof. Suppose $A = UTU^*$, where U is unitary and T is upper triangular. If $\hat{T}$ is the matrix given in Definition 4.2, then by the unitary invariance of the Frobenius norm,

$$\|P_{\mathcal{M}}(U, T) - A\| = \|\hat{T} - T\|.$$

As $\|\hat{T} - T\|^2$ equals the quantity in (4.6), $\|P_{\mathcal{M}}(U, T) - A\|$ depends only on λ and $T_{11}, \dots, T_{nn}$. The result now follows by noting that the T_{ii}'s are the eigenvalues of A and that (4.6) does not depend on the ordering of the T_{ii}'s. $\square$

The next theorem shows that given $A = UTU^*$, if we restrict attention to matrices of the form $U\tilde{T}U^*$, $\tilde{T} \in \mathcal{T}$, then $P_{\mathcal{M}}(U, T)$ is a point in $\mathcal{M}$ closest to A.

THEOREM 4.4. *Suppose $A = UTU^* \in \mathbb{C}^{n \times n}$ with U a unitary matrix and T upper triangular. Then $P_{\mathcal{M}}(U, T)$ satisfies*

$$\|P_{\mathcal{M}}(U, T) - A\| \leq \|U\tilde{T}U^* - A\| \text{ for all } \tilde{T} \in \mathcal{T}.$$

Proof. Let $\tilde{T}$ be a matrix in $\mathcal{T}$. The unitary invariance of the Frobenius norm implies the result will be established if we can show

$$\|\hat{T} - T\| \leq \|\tilde{T} - T\|,$$

where $\hat{T}$ is the matrix given in Definition 4.2. Note that

(4.8) $$\|\tilde{T} - T\|^2 = \sum_{i=1}^{n} |\tilde{T}_{ii} - T_{ii}|^2 + \sum_{i \neq j} |\tilde{T}_{ij} - T_{ij}|^2$$

and that the $\tilde{T}_{ii}$'s are some permutation of the list of eigenvalues λ. The result follows by noting that $\|\hat{T} - T\|^2$ equals the quantity in (4.6) and that this value must be less than or equal to the first summation on the right-hand side of the equality in (4.8).   □

For completeness, we note that, given $A = UTU^*$, $P_{\mathcal{M}}(U, T)$ may not satisfy

$$\|P_{\mathcal{M}}(U, T) - A\| \leq \|M - A\| \text{ for all } M \in \mathcal{M}.$$

For example if

$$U = \frac{1}{5} \begin{bmatrix} -3 & 4 \\ 4 & 3 \end{bmatrix}, \quad T = \begin{bmatrix} 1 & -3 \\ 0 & 2 \end{bmatrix}, \quad \tilde{U} = \frac{1}{5} \begin{bmatrix} -4 & 3 \\ 3 & 4 \end{bmatrix}, \quad \tilde{T} = \begin{bmatrix} 0 & -3 \\ 0 & 0 \end{bmatrix},$$

and $\lambda = \{0, 0\}$, then one can readily verify that

$$\|P_{\mathcal{M}}(U, T) - UTU^*\| \nleq \|\tilde{U}\tilde{T}\tilde{U}^* - UTU^*\|.$$

As for the symmetric case, projection onto $\mathcal{N}$ is straightforward.

THEOREM 4.5. *Given $A \in \mathbb{C}^{n \times n}$, define $A_+ \in \mathbb{R}^{n \times n}$ by*

$$(4.9) \qquad\qquad (A_+)_{ij} = \max\{\mathrm{Re}(A_{ij}), 0\} \text{ for all } 1 \leq i, j \leq n.$$

A_+ is the best approximant in $\mathcal{N}$ to A in the Frobenius norm.

Proof. The projection of $z \in \mathbb{C}$ onto the nonnegative real numbers is given by $\max\{\mathrm{Re}(z), 0\}$. The remainder of the proof follows by exactly the same reasoning used in the proof of Theorem 3.3.   □

Our proposed algorithm for solving the NIEP is the following.

NIEP algorithm:
Problem data. List of desired complex eigenvalues $\lambda = \{\lambda_1, \dots, \lambda_n\}$.
Initialization. Choose a randomly generated nonnegative matrix $Y \in \mathbb{R}^{n \times n}$.
repeat
    1. Calculate a Schur decomposition of Y: $Y = UTU^*$.
    2. $X := P_{\mathcal{M}}(U, T)$.
    3. $Y := X_+$.
  until $\|X - Y\| < \epsilon$.

In the above algorithm, $P_{\mathcal{M}}(U, T)$ is given by Definition 4.2 and X_+ is given by (4.9).

As for the SNIEP algorithm, at each iteration of the NIEP algorithm, X has the desired spectrum λ and Y is nonnegative. If ϵ is small, say $\epsilon = 10^{-14}$, termination of the loop ensures X equals Y (approximately) and hence that Y solves the NIEP.

REMARK 4.6. If each of the members of λ are real and we seek a symmetric nonnegative matrix with spectrum λ, then the NIEP algorithm reduces to the SNIEP algorithm. More precisely, this is true if the members of λ are real, if the initial condition Y is a symmetric nonnegative matrix, and, for Schur decompositions used in the NIEP algorithm, if U is restricted to be real.

Indeed, suppose the current Y is symmetric and nonnegative. For any Schur decomposition of Y, T must be a real diagonal matrix. As we restrict the U matrix to be real, such a decomposition is nothing but a standard eigenvalue-eigenvector

decomposition for a symmetric matrix (though the eigenvalue are not necessarily ordered along the diagonal of T).

As both the elements of λ and the diagonal entries of T are real, the permutation that minimizes (4.6) can be easily characterized. Indeed, in this case (4.6) is minimized if and only if

$$(4.10) \qquad \sum_{i=1}^{n} \hat{\lambda}_i T_{ii}$$

is maximized. From Lemma 3.1, (4.10) is maximized if the $\hat{\lambda}_i$'s are ordered in the same way as the T_{ii}'s. This implies that if Y is symmetric, the step of producing a X from Y is the same in both algorithms.

Lastly, projection of a symmetric matrix onto (4.4) gives the same matrix as projection onto (3.3) and hence this step in both algorithms is also the same. This establishes our claim.

We close this section by noting that unlike the SNIEP algorithm, for the NIEP algorithm there is no guarantee that $\|X - Y\|$ is nonincreasing from one iteration to the next.

5. Convergence. In this section we study the convergence properties of the SNIEP and NIEP algorithms. We present a number of results for the SNIEP algorithm, though limit ourselves to a local convergence result for the NIEP algorithm. We start by characterizing the SNIEP algorithm fixed points. All references to "$\mathcal{M}$," "$\mathcal{N}$," and "the algorithm" refer to the SNIEP versions of these objects, unless otherwise stated.

5.1. Fixed points. As there may be more than one projection of a point Y onto the set $\mathcal{M}$, some care needs to be taken in regard to the definition of fixed points of the algorithm. This subsection includes such a definition, as well as a characterization of these points.

DEFINITION 5.1. $X \in \mathcal{S}^n$ *is a* fixed point *of the algorithm if there exists an eigenvalue-eigenvector decomposition of* X_+,

$$(5.1) \qquad\qquad X_+ = U \tilde{\Lambda} U^T,$$

where $\tilde{\Lambda} = \mathrm{diag}(\tilde{\lambda}_1, \dots, \tilde{\lambda}_n)$, $\tilde{\lambda}_1 \geq \cdots \geq \tilde{\lambda}_n$, *and* U *is orthogonal such that*

$$(5.2) \qquad\qquad X = U \Lambda U^T.$$

X is a fixed point if and only if there is a projection of X_+ onto $\mathcal{M}$ which equals X.

We would like to point out an important fact regarding this definition of a fixed point. If $X \in \mathcal{M}$ is an infeasible fixed point, that is, $X \in \mathcal{M}$ is a fixed point which is not in the solution set $\mathcal{M} \cap \mathcal{N}$, it does *not* necessarily mean that the algorithm cannot make further progress toward a feasible solution from X. This is a consequence of the possible nonuniqueness of the algorithm's matrix decompositions, as we now explain.

If X_+ has distinct eigenvalues, then the orthonormal eigenvectors of X_+ are unique up to multiplication by -1. In this case, any decomposition (5.1) of X_+ will result in the same projected point (5.2). On the other hand, if X_+ has repeated eigenvalues, then there are an infinite number of different decompositions of X_+. If this is the case for an infeasible fixed point X, it may be possible to escape from X by forcing the algorithm to use a different decomposition of X_+.

THEOREM 5.2. *Suppose $X \in \mathcal{M}$ is an infeasible fixed point. Furthermore, suppose that there exists an orthogonal matrix V such that $X_+ = V\tilde{\Lambda}V^T$ and $(V\Lambda V^T)_+ \neq X_+$, that is, that there exists an alternate decomposition of X_+ which leads to a different point in $\mathcal{N}$. Then using this alternate decomposition, the algorithm is able to escape from X.*

Proof. We will show that

$$(5.3) \qquad \|X - X_+\| = \|V\Lambda V^T - X_+\| > \|V\Lambda V^T - (V\Lambda V^T)_+\|.$$

Note that if (5.3) holds, then $V\Lambda V^T \in \mathcal{M}$ and $(V\Lambda V^T)_+ \in \mathcal{N}$ are closer together than X and X_+, and the distance reduction property, Theorem 2.3, implies the result.

The equality in (5.3) holds as both X and $V\Lambda V^T$ are projections of X_+ onto $\mathcal{M}$. The inequality in (5.3) follows by noting that, as $\mathcal{N}$ is closed and convex, $(V\Lambda V^T)_+$ is the unique closest point in $\mathcal{N}$ to $V\Lambda V^T$.   □

For the main result of this subsection, we will need the following lemma.

LEMMA 5.3. *Suppose $x, y \in \mathbb{R}^n$ and $x_1 \geq \cdots \geq x_n$, $y_1 \geq \cdots \geq y_n$. If σ is a permutation of $\{1, \ldots, n\}$ such that*

$$(5.4) \qquad \sum_i x_i y_i = \sum_i x_i y_{\sigma(i)},$$

and, for some i and j,

$$i < j \quad and \quad y_{\sigma(i)} < y_{\sigma(j)},$$

then

$$(5.5) \qquad x_i = x_j.$$

Proof. From Lemma 3.1, for any permutation π of $\{1, \ldots, n\}$,

$$(5.6) \qquad \sum_i x_i y_i \geq \sum_i x_i y_{\pi(i)}.$$

Suppose (5.5) does not hold. Then $x_i > x_j$, which implies

$$(x_i - x_j)(y_{\sigma(i)} - y_{\sigma(j)}) < 0,$$

or rearranging the terms,

$$x_i y_{\sigma(i)} + x_j y_{\sigma(j)} < x_i y_{\sigma(j)} + x_j y_{\sigma(i)}.$$

This combined with (5.4) implies there is a permutation that violates (5.6). As this is not possible, (5.5) must hold.   □

THEOREM 5.4. *Suppose $X \in \mathcal{M}$. X is a fixed point of the algorithm if and only if*

$$(5.7) \qquad [X, X_+] = 0$$

and

$$(5.8) \qquad \operatorname{tr}(\Lambda\tilde{\Lambda}) = \operatorname{tr}(\tilde{\Lambda}^2),$$

where $\tilde{\Lambda} = \mathrm{diag}(\tilde{\lambda}_1, \ldots, \tilde{\lambda}_n)$, $\tilde{\lambda}_1 \geq \cdots \geq \tilde{\lambda}_n$, is the diagonal matrix of eigenvalues of X_+.

Proof. ($\Rightarrow$) If X is a fixed point, there exists an orthogonal matrix U such that

$$(5.9) \qquad X_+ = U\tilde{\Lambda}U^T \quad \text{and} \quad X = U\Lambda U^T.$$

Hence X and X_+ commute and (5.7) holds.

Equality (5.8) follows from

$$\mathrm{tr}(XX_+) = \mathrm{tr}(X_+^2)$$

and (5.9).

($\Leftarrow$) From (5.7), X and X_+ are simultaneously diagonalizable. Hence there exists an orthogonal matrix U and a diagonal matrix $\tilde{\Lambda}_\sigma$ whose diagonal entries are a permutation of the diagonal entries of $\tilde{\Lambda}$ such that

$$X_+ = U\tilde{\Lambda}_\sigma U^T \quad \text{and} \quad X = U\Lambda U^T.$$

By an argument similar to the one used in the first part of the proof,

$$(5.10) \qquad \mathrm{tr}(\Lambda\tilde{\Lambda}_\sigma) = \mathrm{tr}(\tilde{\Lambda}_\sigma^2).$$

Equalities (5.8) and (5.10) imply

$$\mathrm{tr}(\Lambda\tilde{\Lambda}) = \mathrm{tr}(\Lambda\tilde{\Lambda}_\sigma).$$

From Lemma 5.3, if $i < j$ and $(\tilde{\Lambda}_\sigma)_{ii} < (\tilde{\Lambda}_\sigma)_{jj}$, then $\Lambda_{ii} = \Lambda_{jj}$. Hence, the columns of U can always be reordered to get a new U so that (5.9) holds, and hence, X is a fixed point. □

REMARK 5.5. It is interesting to compare the fixed points of the algorithm with those of the SNIEP gradient flow algorithm of [5]. The gradient flow used in [5] is

$$(5.11) \qquad \begin{aligned} \frac{dX}{dt} &= [X, [X, Y]], \\ \frac{dY}{dt} &= 4Y \circ (X - Y). \end{aligned}$$

If $(X(t), Y(t))$, $t \geq 0$, is a solution of this differential equation, then $X(t)$ is isospectral, that is, it preserves the spectrum of $X(0)$, and $Y(t)$ is nonnegative for all $t \geq 0$ if $Y(0)$ is. Suppose $X(0)$ is chosen to have the desired spectrum and $Y(0)$ is chosen nonnegative. Then, if $X(t)$ and $Y(t)$ converge to the same point, that point is a solution of the problem.

The fixed points of (5.11) are the points (X, Y) for which the right-hand side is 0:

$$\begin{aligned} [X, [X, Y]] &= 0, \\ Y \circ (X - Y) &= 0. \end{aligned}$$

Note that for any $X, Y \in \mathcal{S}^n$, $[X, [X, Y]] = 0$ if and only if $[X, Y] = 0$: If $[X, [X, Y]] = 0$, then $0 = \mathrm{tr}(Y[X, [X, Y]]) = \mathrm{tr}([X, Y]^T[X, Y])$ and hence $[X, Y] = 0$. For any $X \in \mathcal{S}^n$, $X_+ \circ (X - X_+) = 0$. Hence, if $X \in \mathcal{M}$ is a fixed point of our SNIEP algorithm, then (X, X_+) is a fixed point of the algorithm of [5]. Roughly speaking,

the set of fixed points of the SNIEP algorithm of this paper is a subset of the set of fixed points of the algorithm of [5].

There do exist infeasible fixed points. If Λ contains negative values (the SNIEP is trivial if it does not), then $X = \Lambda$ is an infeasible fixed point. It may or may not be possible to escape from such a fixed point using alternate decompositions of X_+. An example where escape via this technique is not possible is when $\Lambda = \mathrm{diag}(\lambda_1, \lambda_2)$ with $\lambda_1 > 0$ and $\lambda_2 < 0$, in which case Λ_+ has distinct eigenvalues.

If X is a fixed point, then so is PXP^T for any permutation matrix P. In particular, $P\Lambda P^T$ is a fixed point for any permutation matrix P.

The attractive set of the diagonal fixed points includes the matrices with non-positive off diagonal terms: Suppose X is such a matrix. Then X_+ is diagonal and can be decomposed as $X_+ = P\tilde{\Lambda}P^T$, where $\tilde{\Lambda}$ is diagonal with ordered diagonal entries and P is a permutation matrix. For this decomposition of X_+, Step 2 of the algorithm maps onto the diagonal fixed point $P\Lambda P^T$.

5.2. General convergence properties. The following is a general result regarding convergence to fixed points.

THEOREM 5.6. *Suppose $X_1, X_2, \ldots$, is a sequence of X's produced by the algorithm. Then there is a limit point X of this sequence which is a fixed point of the algorithm and which satisfies*

$$(5.12) \qquad \|X - X_+\| = \lim_{i \to \infty} \|X_i - (X_i)_+\|.$$

If the limit in (5.12) is zero, then every limit point of the sequence is a solution of the problem.

Proof. The X_i's are elements of the compact set $\mathcal{M}$ and hence contain a convergent subsequence. Let X_{i_k} denote the kth element in this subsequence and denote the limit of the X_{i_k}'s by $\hat{X}$. For each k, let U_{i_k} and $\tilde{\Lambda}_{i_k}$ be the matrices from the decomposition of $(X_{i_k})_+$ used to produce X_{i_k+1}, that is, $(X_{i_k})_+ = U_{i_k}\tilde{\Lambda}_{i_k}U_{i_k}^T$ and $X_{i_k+1} = U_{i_k}\Lambda U_{i_k}^T$. As the U_{i_k}'s are members of a compact set, without loss of generality we can assume they converge to a point U. This implies the X_{i_k+1}'s converge, and we denote the corresponding limit point by X.

Corollary 2.4 implies $\lim_{i \to \infty}\|X_i - (X_i)_+\|$ exists. This and the fact that projection onto $\mathcal{N}$ is a continuous operation implies

$$(5.13)$$
$$\|\hat{X} - \hat{X}_+\| = \lim_{k \to \infty} \|X_{i_k} - (X_{i_k})_+\| = \lim_{k \to \infty} \|X_{i_k+1} - (X_{i_k+1})_+\| = \|X - X_+\|.$$

As $(X_{i_k})_+$ converges to $\hat{X}_+$ and the orthogonal matrices U_{i_k} converges to U, it follows that the $\tilde{\Lambda}_{i_k}$'s also converge to, say, $\tilde{\Lambda}$. Hence $\hat{X}_+ = \lim_{k \to \infty}(X_{i_k})_+ = U\tilde{\Lambda}U^T$ and $X = \lim_{k \to \infty} X_{i_k+1} = U\Lambda U^T$. This implies X is a projection of $\hat{X}_+$ onto $\mathcal{M}$.

The equality (5.13) and the fact that X is a projection of $\hat{X}_+$ onto $\mathcal{M}$ imply

$$\|X - X_+\| \geq \|X - \hat{X}_+\|.$$

As X_+ is the unique projection of X onto $\mathcal{N}$, this implies $\hat{X}_+ = X_+$. As X is a projection of $\hat{X}_+ = X_+$ onto $\mathcal{M}$, X is a fixed point.

Suppose now that the limit in (5.12) equals zero. Consider an arbitrary subsequence $X_{i_1}, X_{i_2}, \ldots$, which converges to some point $\tilde{X}$. Note that $\tilde{X}$ is a limit of points in $\mathcal{M}$. The inequality

$$\|(X_{i_k})_+ - \tilde{X}\| \leq \|(X_{i_k})_+ - X_{i_k}\| + \|X_{i_k} - \tilde{X}\|$$

implies it is also a limit of points in $\mathcal{N}$. The last part of the theorem now follows as both $\mathcal{M}$ and $\mathcal{N}$ are closed.   □

The next theorem gives a local convergence result which holds for both SNIEP and NIEP algorithms. If $\mathcal{M}$ and $\mathcal{N}$ are given, respectively, by (3.2) and (3.3), and $\overset{\circ}{\mathcal{N}}$ denotes the interior of $\mathcal{N}$, then, if the intersection of $\mathcal{M}$ and $\overset{\circ}{\mathcal{N}}$ is nonempty, the SNIEP algorithm converges to a solution from points in an open neighborhood of this intersection set. The analogous result for the NIEP algorithm is also true.

THEOREM 5.7. *Let "the algorithm," "the problem," "$\mathcal{M}$" and "$\mathcal{N}$" refer to either "the SNIEP algorithm," "the SNIEP problem," (3.2) and (3.3), or "the NIEP algorithm," "the NIEP problem," (4.3) and (4.4). Suppose $\mathcal{M} \cap \overset{\circ}{\mathcal{N}}$ is nonempty. Then there is an open neighborhood of $\mathcal{M} \cap \overset{\circ}{\mathcal{N}}$ from which the algorithm converges to a solution of the problem (in a single additional iteration).*

Proof. We prove the result for the NIEP algorithm. The proof for the SNIEP algorithm is almost identical.

Suppose $X \in \mathcal{M} \cap \overset{\circ}{\mathcal{N}}$ and let $\epsilon > 0$ be small enough so that the open ball $B(X, \epsilon)$ is a subset of $\mathcal{N}$. Choose $\delta > 0$ such that if $Y \in B(X, \delta)$ and $\tilde{\lambda}_1, \dots, \tilde{\lambda}_n$ are the eigenvalues of Y, then (reordering the $\tilde{\lambda}_i$'s if necessary)

$$(5.14) \qquad \left(\sum_i |\lambda_i - \tilde{\lambda}_i|^2 \right)^{\frac{1}{2}} < \epsilon/2.$$

Here the λ_i's are the desired eigenvalues which define the NIEP. Decreasing δ if necessary, we assume $\delta \leq \epsilon/2$.

We now show that if $Y \in B(X, \delta)$, then any projection of Y onto $\mathcal{M}$ is in $B(X, \epsilon)$. As $B(X, \epsilon) \subset \mathcal{N}$, such a projection of Y is a solution of the problem.

Let $Y \in B(X, \delta)$ and suppose it has Schur decomposition $Y = UTU^*$. Consider

$$\|X - P_{\mathcal{M}}(U, T)\| \leq \|X - Y\| + \|Y - P_{\mathcal{M}}(U, T)\|.$$

The first term on the right of the inequality is less than $\epsilon/2$, as is the second term by the definition of $P_{\mathcal{M}}(U, T)$ and (5.14). This completes the result.   □

As we will see in the next subsection, every feasible $n = 2$ SNIEP has only a finite number of infeasible fixed points. The following result exploits such a situation.

THEOREM 5.8. *Suppose a given feasible SNIEP has only a finite number of infeasible fixed points. There exists a constant $c > 0$ such that if $X_1, X_2, \dots$, is a sequence of X's produced by the algorithm and $\|X_{i^*} - (X_{i^*})_+\| < c$ for some i^*, then any limit point of the X_i's (there must be at least one) is a solution of the problem.*

Proof. Let $c > 0$ be such that $\|Z - Z_+\| \geq c$ for all infeasible fixed points Z. Such a c exists as there are only a finite number of infeasible fixed points. By Corollary 2.4, $\|X_i - (X_i)_+\|$ is a nonincreasing function of i and hence must have a limit. Theorem 5.6 implies this limit must be zero. The rest now also follows from Theorem 5.6.   □

If a limit point in Theorem 5.8 is in $\mathcal{M} \cap \overset{\circ}{\mathcal{N}}$, then Theorem 5.7 implies the sequence of X_i's will converge to a solution (in a finite number of iterations).

Do all feasible SNIEPs have only a finite number of infeasible fixed points? Alternatively, as $\mathcal{M}$ is compact, an equivalent question is: Are all infeasible fixed points of a feasible SNIEP isolated? These are interesting questions, to which we currently do not have an answer.

5.3. Further analysis: $n = 2$ SNIEP. In this subsection we continue our analysis of convergence; in particular we investigate the $n = 2$ SNIEP. Though necessary and sufficient conditions exist for the solvability of the $n = 2$ SNIEP, and there exists an analytic solution when these conditions are met, we believe the analysis presented here is still quite interesting and also gives insight into higher-dimensional problems.

As noted in [5], for $n = 2$, feasible SNIEPs have a very nice geometric interpretation. If the 2×2 symmetric matrices are parameterized by $\mathbb{R}^3$ in the standard way and the eigenvalues defining $\mathcal{M}$ are distinct, then the points with the desired spectrum form a one dimensional ellipse in $\mathbb{R}^3$, and the SNIEP is equivalent to finding a point on this ellipse that is also in the nonnegative orthant of $\mathbb{R}^3$.

As the trace of a matrix equals the sum of its eigenvalues, a necessary condition for solvability is that $\lambda_1 + \lambda_2 \geq 0$. In fact, this condition is also sufficient; if it is met, then a solution of the problem is

$$(5.15) \qquad X = \frac{1}{2} \begin{pmatrix} \lambda_1 + \lambda_2 & \lambda_1 - \lambda_2 \\ \lambda_1 - \lambda_2 & \lambda_1 + \lambda_2 \end{pmatrix}.$$

As normal, here we assume the eigenvalues are ordered: $\lambda_1 \geq \lambda_2$.

The feasible cases can be enumerated as follows:
1. $\lambda_1 = \lambda_2 \geq 0$,
2. $\lambda_1 > \lambda_2 \geq 0$,
3. $\lambda_1 > 0 > \lambda_2$, $\lambda_1 \geq |\lambda_2|$.

Theorem 5.10 below characterizes the infeasible fixed points of the algorithm for the different cases listed above. We will need the following lemma.

LEMMA 5.9. *Suppose*

$$X = \begin{pmatrix} a & b \\ b & a \end{pmatrix}$$

for some a and some $b \leq 0$. Then $X \in \mathcal{M}$ and is a fixed point if and only if

$$(5.16) \qquad X = \frac{1}{2} \begin{pmatrix} \lambda_1 + \lambda_2 & \lambda_2 - \lambda_1 \\ \lambda_2 - \lambda_1 & \lambda_1 + \lambda_2 \end{pmatrix}.$$

Proof. ($\Rightarrow$) $X \in \mathcal{M}$ and $b \leq 0$ implies $\lambda_1 = a - b$ and $\lambda_2 = a + b$. Solving for a and b gives (5.16).
($\Leftarrow$) If X is given by (5.16), then its eigenvalues are λ_1 and λ_2, and hence it is a member of $\mathcal{M}$. As $\lambda_2 - \lambda_1 \leq 0$, X_+ is a constant multiple of the identity. If U is any orthogonal matrix such that $X = U\Lambda U^T$, then $X_+ = U\tilde{\Lambda}U^T$ with $\tilde{\Lambda} = X_+$ and hence X is a fixed point. □

THEOREM 5.10. *For case 1, $\mathcal{M}$ contains only the single element (5.15) (which equals $\lambda_1 I$) and hence in this case the algorithm does not have any infeasible fixed points. For case 2, there is a single infeasible fixed point, (5.16). For case 3, there are three infeasible fixed points, (5.16),*

$$(5.17) \qquad X = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix},$$

and

$$(5.18) \qquad X = \begin{pmatrix} \lambda_2 & 0 \\ 0 & \lambda_1 \end{pmatrix}.$$

Proof. We assume $X \in \mathcal{M}$ is a fixed point given by

$$X = \begin{pmatrix} a & b \\ b & c \end{pmatrix}$$

and consider all possibilities for a, b, and c.

Case 2. If $a = 0$ or $c = 0$, then $\det(X) \geq 0$ implies $b = 0$, which implies X is feasible. If $a > 0 > c$ or $c > 0 > a$, then $\det(X) < 0$, which is not possible. If $a < 0$ and $c < 0$, then $\mathrm{tr}(X) < 0$, which is not possible. Hence it remains to consider the subcase $a > 0$ and $c > 0$.

Suppose $a > 0$ and $c > 0$. If $b \geq 0$, X is feasible. If $b < 0$, then $X_+ = \mathrm{diag}(a, c)$. If $a \neq c$, then the orthonormal eigenvectors of X_+, up to multiplication by -1, are the standard orthonormal basis vectors for $\mathbb{R}^2$. This would imply X is diagonal, contradicting $b < 0$. If $a = c$, by Lemma 5.9, the only $X \in \mathcal{M}$ which is a fixed point is given by (5.16). As $\lambda_2 - \lambda_1 < 0$, X is infeasible.

Case 3. By considering the equalities $\mathrm{tr}(X) = \lambda_1 + \lambda_2$ and $\det(X) = \lambda_1\lambda_2$, it follows that

$$X = \begin{pmatrix} a & b \\ b & \lambda_1 + \lambda_2 - a \end{pmatrix}, \quad b = \pm\sqrt{(\lambda_1 - a)(a - \lambda_2)}, \quad \text{and that } a \in [\lambda_2, \lambda_1].$$

Suppose $b \leq 0$. Then $X_+ = \mathrm{diag}(a_+, (\lambda_1+\lambda_2-a)_+)$. X_+ has repeated eigenvalues if and only if $a = (\lambda_1+\lambda_2)/2$. If X_+ does have repeated eigenvalues, then the diagonal terms of X are equal and by Lemma 5.9 the only $X \in \mathcal{M}$ which is a fixed point is given by (5.16). As $\lambda_2 - \lambda_1 < 0$, X is infeasible. If X_+ has distinct eigenvalues, then X must be one of the infeasible fixed points (5.17) or (5.18).

Suppose $b > 0$. If $a \in [0, \lambda_1 + \lambda_2]$, X is feasible. It is not possible that $a = \lambda_1$ or $a = \lambda_2$ as then, $b = 0$. Hence it remains to consider $a \in (\lambda_1 + \lambda_2, \lambda_1)$ (the case $a \in (\lambda_2, 0)$ follows from this case by replacing a with $\lambda_1 + \lambda_2 - a$). If X_+ and X are given by (5.1) and (5.2), respectively, then

$$X - X_+ = \begin{pmatrix} 0 & 0 \\ 0 & \lambda_1 + \lambda_2 - a \end{pmatrix} = U(\Lambda - \tilde{\Lambda})U^T.$$

This implies $\Lambda - \tilde{\Lambda}$ has distinct eigenvalues and hence that (up to multiplication of its columns by -1)

$$U = I \quad \text{or} \quad U = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

This implies X is diagonal but this contradicts the fact that $b > 0$ and hence this subcase cannot occur.   □

Consider the infeasible fixed point X given by (5.16). The U satisfying (5.1) and (5.2) is unique up to multiplication of its columns by -1. It is given by

$$U = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}.$$

Notice that X_+ is a constant multiple of the identity. If $X_+ = V\tilde{\Lambda}V^T$ is any alternate decomposition of X_+, that is, if V is *any* orthogonal matrix that does not equal U nor U with one or both of its columns multiplied by -1, then $V\Lambda V^T \neq X$, and

Theorem 5.10 implies $(V\Lambda V^T)_+ \neq X_+$. Hence, Theorem 5.2 implies, using almost any decomposition of X_+, the algorithm is able to escape from X.

The only other infeasible fixed points are the diagonal ones given by (5.17) and (5.18). For these fixed points, X_+ has distinct eigenvalues, and alternate decompositions as a means of escape cannot be utilized. Despite this, the next theorem shows that such fixed points are unstable and one can escape from them by adding an arbitrarily small perturbation.

THEOREM 5.11. *Suppose X is the fixed point (5.17) or (5.18) ($\lambda_1 > 0 > \lambda_2$). Then there exist a positive constant $\bar{b}$ such that for each b satisfying $0 < b \leq \bar{b}$ there exist positive $\bar{a} = \bar{a}(b)$ and $\bar{c} = \bar{c}(b)$ such that if $|a| \leq \bar{a}$ and $|c| \leq \bar{c}$, then replacing X by $X + P$, where*

$$P = \begin{pmatrix} a & b \\ b & c \end{pmatrix},$$

leads to escape from the fixed point X.

Proof. To prove the theorem we will show that if P is as above and $(X + P)_+ = V\tilde{\Lambda}V^T$, then

(5.19) $\|X_+ - X\| > \|(X + P)_+ - V\Lambda V^T\|.$

Inequality (5.19) and Theorem 2.3 together imply we cannot return to X.

We assume the fixed point X is given by (5.17). (The proof of the (5.18) case is identical except for a permutation of matrix rows and columns.)

Suppose $a \geq -\lambda_1$, $b > 0$, and $c \leq -\lambda_2$. Then

(5.20) $(X + P)_+ = \begin{pmatrix} \lambda_1 + a & b \\ b & 0 \end{pmatrix}.$

If $\tilde{\Lambda} = \mathrm{diag}(\tilde{\lambda}_1, \tilde{\lambda}_2)$, then

$$\tilde{\lambda}_1, \tilde{\lambda}_2 = \frac{\lambda_1 + a \pm \sqrt{(\lambda_1 + a)^2 + 4b^2}}{2}.$$

By the unitary invariance of the Frobenius norm, (5.19) is equivalent to

$$\lambda_2^2 > (\lambda_1 - \tilde{\lambda}_1)^2 + (\lambda_2 - \tilde{\lambda}_2)^2.$$

Substituting $\tilde{\lambda}_1 + \tilde{\lambda}_2 = \lambda_1 + a$, we have

(5.21) $\lambda_2^2 > (\tilde{\lambda}_2 - a)^2 + (\lambda_2 - \tilde{\lambda}_2)^2.$

For now, suppose $a = 0$. Noting that $\tilde{\lambda}_2 < 0$ as $b \neq 0$, straightforward algebraic manipulations imply (5.21) is equivalent to $\lambda_2 < \tilde{\lambda}_2$. Hence, for $a = 0$, (5.19) holds if and only if $\lambda_1 - \sqrt{\lambda_1^2 + 4b^2} > 2\lambda_2$. As $\lambda_1 > 0 > \lambda_2$, this inequality holds for all $b > 0$ small enough.

Each $b > 0$ that satisfies (5.21) when $a = 0$ also satisfies this inequality for all a sufficiently small as the right-hand side of (5.21) depends continuously on a.   □

Suppose that for $i = 1, 2, \ldots$, $(a, b) = (a_i, b_i)$, $b_i > 0$, satisfies (5.21) and that the b_i's converge to zero. Examination of (5.21) and the expression for $\tilde{\lambda}_2$ shows that the a_i's must also converge to zero, and hence this theorem is the best we can do.

Theorem 5.11 can be readily extended to infeasible diagonal fixed points of problems of size $n > 2$.

SNIEP: A comparison of performance for different problem sizes n. i denotes the average number of iterations and T denotes the average convergence time in CPU seconds.

n	i	T	% solved
5	19	0.0016	100
10	18	0.0030	100
20	17	0.0075	100
100	12	0.15	100

6. Numerical experiments. This section contains some numerical results for both the SNIEP and NIEP algorithms.

All computational results were obtained using a 3 GHz Pentium 4 machine. The algorithms were coded using Matlab 7.0.

Throughout this section, when we say a matrix is "randomly generated" we mean each entry of that matrix is randomly drawn from the uniform distribution on the interval $[0, 1]$. When dealing with the SNIEP algorithm, all randomly generated matrices are chosen symmetric.

For both algorithms, the initial starting Y is always randomly generated and the convergence tolerance ϵ is set to 10^{-14}.

A final note before presenting the results: Suppose $\mathcal{M} \cap \overset{\circ}{\mathcal{N}}$ is nonempty and X is a member of this set. Then for any real orthogonal matrix Q that is sufficiently close to the identity, QXQ^T is also a solution. In particular, if $\mathcal{M} \cap \overset{\circ}{\mathcal{N}}$ is nonempty, then there will be multiple solutions. This comment applies to both SNIEPs and NIEPs.

6.1. SNIEP. This subsection starts with some results for randomly generated SNIEPs. To ensure each problem is feasible, each desired spectrum is taken from a randomly generated matrix.

Results for various problem sizes n are given in Table 6.1. For each value of n, 1000 problems were considered. The table contains the average number of iterations required to find a solution, the average time required to find a solution, and the success rate. As can be seen, the algorithm performed extremely well and was able to solve every problem. In all cases, both the average number of iterations and the average solution time was very small.

REMARK 6.1. It is interesting to note that T increases with n, as would be expected, while i decreases. A reason for this could be the following. As already mentioned, for any choice of desired eigenvalues, $\mathcal{M}$ is a smooth manifold. In addition, if the eigenvalues defining $\mathcal{M}$ are distinct, as they will be if they were taken from a randomly generated matrix, then the dimension of $\mathcal{M}$ is $n(n-1)/2$; see [14, Chapter 2]. The dimension of $\mathcal{S}^n$ is $n(n+1)/2$. Hence,

$$\frac{\dim \mathcal{M}}{\dim \mathcal{S}^n} = \frac{n-1}{n+1},$$

which is an increasing function of n. For larger n, $\mathcal{M}$ is "thicker" relative to the ambient space and hence, intuitively, the corresponding SNIEP is easier to solve.

Suppose $X_1, X_2, \dots,$ is a sequence of X's produced by the SNIEP algorithm and that these points converge to a solution $\bar{X}$. Figure 6.1 shows a typical plot of $\|X_i - \bar{X}\|$ versus i. Convergence is clearly linear. This is to be expected: Suppose $\bar{X}$ is a point on the boundary of $\mathcal{N}$ and that the $(X_i)_+$'s lie in a particular face of $\mathcal{N}$. As $\mathcal{M}$ is a manifold, near $\bar{X}$ it looks locally like an affine subspace of $\mathcal{S}^n$. As the face of $\mathcal{N}$

TABLE 6.2

SNIEP: A problem with repeated eigenvalues, $\lambda = \{3 - t, 1 + t, -1, -1, -1, -1\}$. i denotes the average number of iterations and T denotes the average convergence time in CPU seconds. i and T do not include the attempts that had not converged after 5000 iterations.

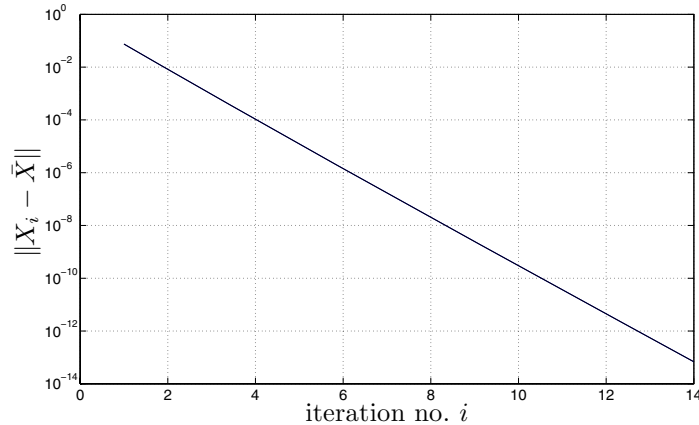| t	i	T	% solved
0.25	480	0.061	100
0.5	470	0.061	97
0.75	340	0.050	65
0.95	310	0.046	59

FIG. 6.1. *Linear convergence of the SNIEP algorithm.*

also looks locally like an affine subspace, we would expect local linear convergence as alternating projections between two intersecting affine subspaces converge linearly [11].

Randomly generated problems have properties that are not shared by all SNIEPs. For example, as already mentioned, randomly generated problems have distinct eigenvalues. We next consider a problem with repeated eigenvalues, namely $\lambda = \{3 - t, 1 + t, -1, -1, -1, -1\}$ for $0 < t < 1$. The $t = 1/2$ version of the problem is also considered in [5], where a numerical solution is sought via the gradient flow approach of that paper. An analytic solution to this problem is given in [20].

Notice that for any value of t the desired eigenvalues sum to zero and hence there exist arbitrarily small perturbations of the spectrum which lead to an infeasible SNIEP. In particular this problem cannot have any solutions in the interior of $\mathcal{N}$. We have tried the SNIEP algorithm on a number of other problems with repeated eigenvalues with excellent results. This is the hardest problem we have encountered so far.

The results of applying the algorithm to the problem for various values of t are given in Table 6.2. They are based on running the algorithm 100 times for each value of t.

First, the results indicate that the SNIEP algorithm is not always successful in finding a solution. However, they also show that the algorithm can still be quite successful if a number of initial conditions are tried. It is interesting to note that the algorithm becomes more sensitive to the choice of the initial condition the larger t is. Notice that as $t \to 1$, the eigenvalues $3 - t$ and $1 + t$ both converge to the same value,

TABLE 6.3
NIEP: A comparison of performance for different problem sizes n. i denotes the average number of iterations and T denotes the average convergence time in CPU seconds. i and T do not include the problems that had not converged after 5000 iterations.

n	i	T	% solved
5	26	0.011	99.7
10	44	0.045	99.8
20	48	0.12	99.8
100	200	12	96.6

and the dimension of the manifold $\mathcal{M}$ (which depends solely on the multiplicities of the eigenvalues) goes from 9 when $0 < t < 1$ to 8 when $t = 1$ [14, Chapter 2].

Aside: Regarding initial conditions, as noted before, both the SNIEP and NIEP algorithms use a nonnegative initial starting point. This is important and, in fact, the performance of neither algorithm is as good if non-nonnegative initial conditions are used.

Here is a solution that was found to the $t = 1/2$ problem:

$$X = \begin{pmatrix} 0 & 0 & 0 & \sqrt{\frac{3}{2}} & 0 & 1 \\ 0 & 0 & 1 & \frac{1}{2} & 1 & 0 \\ 0 & 1 & 0 & \frac{1}{2} & 1 & 0 \\ \sqrt{\frac{3}{2}} & \frac{1}{2} & \frac{1}{2} & 0 & \frac{1}{2} & \sqrt{\frac{3}{2}} \\ 0 & 1 & 1 & \frac{1}{2} & 0 & 0 \\ 1 & 0 & 0 & \sqrt{\frac{3}{2}} & 0 & 0 \end{pmatrix}.$$

This solution is different to both the solution in [20] and the solution in [5]. A number of other solutions were also found.

Here is a X_+ corresponding to an infeasible X (again for the $t = 1/2$ problem):

$$X_+ = \begin{pmatrix} \frac{1}{2} & 0 & 0 & 0 & 0 & \sqrt{\frac{3}{2}} \\ 0 & 0 & \frac{7}{8} & \frac{7}{8} & \frac{7}{8} & 0 \\ 0 & \frac{7}{8} & 0 & \frac{7}{8} & \frac{7}{8} & 0 \\ 0 & \frac{7}{8} & \frac{7}{8} & 0 & \frac{7}{8} & 0 \\ 0 & \frac{7}{8} & \frac{7}{8} & \frac{7}{8} & 0 & 0 \\ \sqrt{\frac{3}{2}} & 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

The eigenvalues of this matrix are $\tilde{\lambda} = \{2\frac{5}{8}, 1\frac{1}{2}, -\frac{7}{8}, -\frac{7}{8}, -\frac{7}{8}, -1\}$.

6.2. NIEP. This subsection starts with some results for randomly generated NIEPs. Again, to ensure each problem is feasible, each desired spectrum is taken from a randomly generated matrix. Results are given in Table 6.3.

As can be seen, the results are again very good, with almost all problems solved.

The results indicate that NIEPs are harder to solve than SNIEPs. Also, the number of iterations, time, and time per iteration are greater. Part of the reason for an increase in time per iteration will be the extra computation required to calculate the least squares matching component of each $P_{\mathcal{M}}(U, T)$ calculation; see (4.6). (For

SNIEPs, the corresponding step is easy: the eigenvalues are real and just need to be sorted in decreasing order.)

For the NIEPs, both i and T increased with n.

The final problem we consider is taken from [8]. It is to find a stochastic matrix with (presumedly randomly generated) spectrum $\lambda = \{1.0000, -0.2608, 0.5046, 0.6438, -0.4483\}$. Furthermore the problem requires the matrix to have zeros in certain positions. In the context of Markov chains, we require the states to form a ring and that each state be linked to at most two immediate neighbors. The zero pattern is given by the zeros of the following matrix:

$$
(6.1) \qquad Z = \begin{pmatrix} 1 & 1 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 \end{pmatrix}.
$$

Our algorithm as it stands is not able to solve this problem though it is able to do so if a simple modification is made. Using Z from (6.1), define

$$
\tilde{\mathcal{N}} = \{A \in \mathbb{R}^{n \times n} \mid A_{ij} \geq 0 \text{ and } A_{ij} = 0 \text{ if } Z_{ij} = 0\}.
$$

$\tilde{\mathcal{N}}$ is still a convex set. In the NIEP algorithm, replacing projection onto $\mathcal{N}$ by projection onto $\tilde{\mathcal{N}}$ gives solutions (nonnegative matrices) with zeros in the desired places. Using the transformation discussed in the introduction of the paper, solutions found by the algorithm can be converted into stochastic matrices with the same spectrum. Note that this transformation preserves zeros.

Using this methodology readily produced many solutions. An example is

$$
X = \begin{pmatrix} 0.6931 & 0.2887 & 0 & 0 & 0.0182 \\ 0.1849 & 0.2422 & 0.5729 & 0 & 0 \\ 0 & 0.5476 & 0.3622 & 0.0902 & 0 \\ 0 & 0 & 0.5437 & 0.1233 & 0.3330 \\ 0.3712 & 0 & 0 & 0.6103 & 0.0185 \end{pmatrix}.
$$

Another solution is

$$
X = \begin{pmatrix} 0.8634 & 0.0431 & 0 & 0 & 0.0936 \\ 0.6224 & 0 & 0.3776 & 0 & 0 \\ 0 & 0.4935 & 0.1564 & 0.3501 & 0 \\ 0 & 0 & 0.1107 & 0.0115 & 0.8778 \\ 0.3452 & 0 & 0 & 0.2467 & 0.4080 \end{pmatrix}.
$$

Notice that this latter solution has an extra zero. While this X still solves the problem, by further modifying $\mathcal{N}$ it is possible to ensure zeros appear only in the places specified by (6.1) and nowhere else.

For example, using

$$
\tilde{\mathcal{N}} = \{A \in \mathbb{R}^{n \times n} \mid A_{ij} = 0 \text{ if } Z_{ij} = 0 \text{ and } A_{ij} \geq \delta \text{ otherwise}\},
$$

with $\delta > 0$ a small constant, does the trick. Note that the stochastic matrix transformation leaves positive entries positive.

7. Conclusion. In this paper we have presented two related numerical methods, one for the NIEP, which can also be used to solve the inverse eigenvalue problem for stochastic matrices, and another for the SNIEP. The ideas used in the paper should also be applicable to many other inverse eigenvalue problems, including other problems involving nonsymmetric matrices.

8. Appendix. Local uniqueness and smoothness of projections. This appendix contains some supplementary results regarding projection onto the symmetric version of $\mathcal{M}$; see (3.2). While these results are not used in the main body of the paper, we believe they are interesting and worth mentioning. We would also expect them to be useful for other inverse eigenvalue problems.

THEOREM 8.1. *There is an open neighborhood of $\mathcal{M}$ from which projections onto $\mathcal{M}$ are unique.*

Proof. To ease the presentation we will assume Λ has only two distinct eigenvalues. (The general case follows by similar reasoning.) Let m be such that $\lambda_m > \lambda_{m+1}$. If $Z \in \mathcal{M}$, then there exists a neighborhood of Z such that each matrix in this neighborhood has distinct mth and $m+1$th (ordered) eigenvalues. Suppose Y is an element in this neighborhood with eigenvalues $\tilde{\lambda}_1 \geq \cdots \geq \tilde{\lambda}_n$. It follows from the proof of Theorem 3.2 that if X is a projection of Y onto $\mathcal{M}$, then there exist orthonormal vectors $u_1, \ldots, u_n$ such that $Y = \sum_{i=1}^{n} \tilde{\lambda}_i u_i u_i^T$ and $X = \lambda_1 \sum_{i=1}^{m} u_i u_i^T + \lambda_{m+1} \sum_{i=m+1}^{n} u_i u_i^T$. Note that here we have used the fact that $\tilde{\lambda}_m$ and $\tilde{\lambda}_{m+1}$ are distinct. The proof will be complete if we can show X does not depend on the particular decomposition of Y.

Suppose eigenvalue $\tilde{\lambda}_j$ has multiplicity k with $\tilde{\lambda}_j = \cdots = \tilde{\lambda}_{j+k-1}$. If $\hat{u}_j, \ldots, \hat{u}_{j+k-1}$ is another set of orthonormal vectors that span the eigenspace corresponding to $\tilde{\lambda}_j$, then there exists a $k \times k$ orthogonal matrix Θ such that $[\hat{u}_j, \ldots, \hat{u}_{j+k-1}] = [u_j, \ldots, u_{j+k-1}]\Theta$. Consequently,

$$\sum_{i=j}^{j+k-1} \hat{u}_i \hat{u}_i^T = \sum_{i=j}^{j+k-1} u_i u_i^T.$$

The separation of eigenvalues implies the indices $j, \ldots, j+k-1$ are all either less than or equal to m, or, greater than or equal to $m+1$. It follows that X does not depend on the particular decomposition of Y. ☐

THEOREM 8.2. *There is an open neighborhood of $\mathcal{M}$ from which the (unique) projections onto $\mathcal{M}$ are smooth.*

Proof (Outline). $\mathcal{M}$ is a submanifold of $\mathcal{S}^n$ and hence each point in $\mathcal{M}$ is in the image of a local parametrization of $\mathcal{M}$. The result can be shown to hold locally by using such a parametrization, using a condition necessary for a point to be a projection (if X is a projection of Y, then $X - Y$ is normal to the tangent space of $\mathcal{M}$ at X), and employing the implicit function theorem. In trying to satisfy the conditions of the implicit function theorem, the requirement that points being projected are sufficiently close to $\mathcal{M}$ appears.

As a consequence of the above mentioned necessary condition not being sufficient, for the proof to work it appears to be that it must be known a priori that in a neighborhood with unique projections, the projection operation is continuous. This is indeed the case as can be shown via a contradiction argument. ☐

The proof of the above theorem does not use any properties of $\mathcal{M}$ aside from the fact that it is a (closed) submanifold and that, near the set, projections are unique. (The projection result that uniqueness implies continuity holds for projections onto any closed set.) Hence, Theorem 8.2 also holds for any set with these two properties.

REFERENCES

[1] H. H. Bauschke and J. M. Borwein, *On projection algorithms for solving convex feasibility problems*, SIAM Rev., 38 (1996), pp. 367–426.

[2] A. Berman and R. J. Plemmons, *Nonnegative Matrices in the Mathematical Sciences*, Academic Press, New York, 1979. Also published as Classics in Appl. Math. 9, SIAM, Philadelphia, 1994.

[3] L. M. Brègman, *The method of successive projection for finding a common point of convex sets*, Soviet Mathematics, 6 (1965), pp. 688–692.

[4] M. T. Chu, *Inverse eigenvalue problems*, SIAM Rev., 40 (1998), pp. 1–39.

[5] M. T. Chu and K. R. Driessel, *Constructing symmetric nonnegative matrices with prescribed eigenvalues by differential equations*, SIAM J. Math. Anal., 22 (1991), pp. 1372–1387.

[6] M. T. Chu and G. H. Golub, *Structured inverse eigenvalue problems*, Acta Numer., 11 (2002), pp. 1–71.

[7] M. T. Chu and G. H. Golub, *Inverse Eigenvalue Problems: Theory, Algorithms, and Applications*, Oxford University Press, Oxford, 2005.

[8] M. T. Chu and Q. Guo, *A numerical method for the inverse stochastic spectrum problem*, SIAM J. Matrix Anal. Appl., 19 (1998), pp. 1027–1039.

[9] M. T. Chu and S. F. Xu, *On computing minimal realizable spectral radii of non-negative matrices*, Numer. Linear Algebra Appl., 12 (2005), pp. 77–86.

[10] P. L. Combettes and H. J. Trussell, *Method of successive projections for finding a common point of sets in metric spaces*, J. Optim. Theory Appl., 67 (1990), pp. 487–507.

[11] F. Deutsch, *Best Approximation in Inner Product Spaces*, Springer-Verlag, New York, 2001.

[12] P. D. Egleston, T. D. Lenker, and S. K. Narayan, *The nonnegative inverse eigenvalue problem*, Linear Algebra Appl., 379 (2004), pp. 475–490.

[13] G. H. Hardy, J. E. Littlewood, and G. Polya, *Inequalities*, Cambridge University Press, Cambridge, 1952.

[14] U. Helmke and J. B. Moore, *Optimization and Dynamical Systems*, Springer-Verlag, London, 1994.

[15] R. A. Horn and C. R. Johnson, *Matrix Analysis*, Cambridge University Press, Cambridge, 1985.

[16] C. R. Johnson, T. J. Laffey, and R. Loewy, *The real and the symmetric nonnegative inverse eigenvalue problems are different*, Proc. Amer. Math. Soc., 124 (1996), pp. 3647–3651.

[17] D. Luenberger, *Optimization by Vector Space Methods*, Wiley, New York, 1969.

[18] H. Minc, *Nonnegative matrices*, Wiley, New York, 1988.

[19] H. Perfect, *Methods of constructing certain stochastic matrices*, Duke Math. J., 20 (1953), pp. 395–404.

[20] G. W. Soules, *Constructing symmetric nonnegative matrices*, Linear and Multilinear Algebra, 13 (1983), pp. 241–251.

[21] K. R. Suleimanova, *Stochastic matrices with real characteristic numbers*, Soviet Math. Dokl., 66 (1949), pp. 343–345. (In Russian).

[22] S. F. Xu, *An Introduction to Inverse Algebraic Eigenvalue Problems*, Peking University Press, Beijing, and Vieweg & Sohn, Braunschweig, Germany, 1998.

[23] K. Yang and R. Orsi, *Pole placement via output feedback: a methodology based on projections*, in Proceedings of the 16th IFAC World Congress, Prague, Czech Republic, 2005.

[24] K. Yang, R. Orsi, and J. B. Moore, *A projective algorithm for static output feedback stabilization*, in Proceedings of the 2nd IFAC Symposium on System, Structure and Control, Oaxaca, Mexico, 2004, pp. 263–268.