A geometric Newton method for Oja's vector field*

P.-A. Absil† M. Ishteva‡ L. De Lathauwer§‡ S. Van Huffel‡

June 21, 2008

Abstract

Newton's method for solving the matrix equation $F(X) \equiv AX - XX^TAX = 0$ runs up against the fact that its zeros are not isolated. This is due to a symmetry of F by the action of the orthogonal group. We show how differential-geometric techniques can be exploited to remove this symmetry and obtain a "geometric" Newton algorithm that finds the zeros of F. The geometric Newton method does not suffer from the degeneracy issue that stands in the way of the original Newton method.

Key words. Oja's learning equation, Oja's flow, differential-geometric optimization, Riemannian optimization, quotient manifold, neural networks

*This paper presents research results of the Belgian Network DYSCO (Dynamical Systems, Control, and Optimization), funded by the Interuniversity Attraction Poles Programme, initiated by the Belgian State, Science Policy Office. The scientific responsibility rests with its authors. Research supported by: (1) Research Council K.U.Leuven: GOA-Ambiorics, CoE EF/05/006 Optimization in Engineering (OPTEC), CIF1, (2) F.W.O.: (a) project G.0321.06, (b) Research Communities ICCoS, ANMMM and MLDM, (3) the Belgian Federal Science Policy Office: IUAP P6/04 (DYSCO, "Dynamical systems, control and optimization", 2007–2011), (4) EU: ERNSI. M. Ishteva is supported by a K.U.Leuven doctoral scholarship (OE/06/25, OE/07/17), L. De Lathauwer is supported by "Impulsfinanciering Campus Kortrijk (2007-2012)(CIF1)".

†Département d'ingénierie mathématique, Université catholique de Louvain, Belgium (www.inma.ucl.ac.be/∼absil).

‡Research Division SCD of the Department of Electrical Engineering (ESAT) of the Katholieke Universiteit Leuven, Kasteelpark Arenberg 10, B-3001 Leuven, Belgium.

§Subfaculty Science and Technology of the Katholieke Universiteit Leuven Campus Kortrijk, E. Sabbelaan 53, 8500 Kortrijk, Belgium.

1 Introduction

Let A be a symmetric positive-definite $n \times n$ matrix and let p be a positive integer smaller than n.

Oja's flow (Oja, 1982, 1989), in its averaged version

$$\frac{\mathrm{d}}{\mathrm{d}t}X(t) = F(X(t)), \tag{1a}$$

where F denotes the vector field

$$F : \mathbb{R}^{n\times p} \to \mathbb{R}^{n\times p} : X \mapsto AX - XX^TAX, \tag{1b}$$

is well known for its principal component analysis properties. For all initial conditions X_0, the ordinary differential equation (1) has a unique solution curve $t \mapsto X(t)$ on the interval $[0, \infty)$ (Yan, Helmke, & Moore, 1994, Th. 2.1), the limit $X(\infty) = \lim_{t\to\infty} X(t)$ exists, the convergence to $X(\infty)$ is exponential, and $X(\infty)$ is a zero of Oja's vector field F (1b) (Yan et al., 1994, Th. 3.1).

Observe that the zeros of F are the solutions $X \in \mathbb{R}^{n\times p}$ of the matrix equation

$$AX = XX^TAX,$$

which implies that the columns of AX are linear combinations of the columns of X. Letting

$$\mathrm{col}(Y) = \{Y\alpha : \alpha \in \mathbb{R}^p\}$$

denote the column space of $Y \in \mathbb{R}^{n\times p}$, it follows that all zeros X of (1b) satisfy $A\,\mathrm{col}(X) \subseteq \mathrm{col}(X)$, i.e., $\mathrm{col}(X)$ is an invariant subspace of A. Moreover, $X(\infty)$ is an orthonormal matrix ($X^T(\infty)X(\infty) = I_p$), thus of full rank, whenever the initial condition $X(0)$ is of full rank (Yan et al., 1994, Prop. 3.1). Assuming that $X(0)$ has full rank, it holds that $\mathrm{col}(X(\infty))$ is the p-dimensional principal subspace of A (assumed to be unique) if and only if $\mathrm{rank}(X_0^TV_p) = p$, where V_p is any $n\times p$ matrix whose column space is the p-dimensional principal subspace of A (Yan et al., 1994, Th. 5.1). (This condition means that $\mathrm{col}(X(0))$ does not contain any direction orthogonal to the p-dimensional principal subspace of A.) The set of all initial conditions that do not satisfy this condition is a negligible set, i.e., Oja's flow asymptotically computes the p-dimensional principal subspace of A for almost all initial conditions. We also point out that Oja's flow induces a subspace flow, called the power flow (Absil, Sepulchre, & Mahony, 2008).

Because of these remarkable properties, Oja's flow has been and remains the subject of much attention, in its form (1) as well as in several modified versions; see, for example, (Yan et al., 1994; Yan, 1998; Dehaene, Moonen, & Vandewalle, 1999; Manton, Helmke, & Mareels, 2005; Jankovic & Ogawa, 2006) . However, turning Oja's flow into a principal component algorithm implementable in a digital computer requires to discretize the flow in such a way that its good convergence properties are preserved. Since the solutions $X(t)$ converge exponentially to their limit point $X(\infty)$, it follows that the sequence of equally spaced discrete-time samples $(X(kT))_{k\in\mathbb{N}}$ converges only Q-linearly (Ortega & Rheinboldt, 1970) to $X(\infty)$. Therefore, numerical integration methods that try to compute accurately the solution trajectories of Oja's flow are not expected to converge faster than linearly.

Nevertheless, if the ultimate goal is to compute the principal eigenspace of A, then it is tempting to try to accelerate the convergence when the iterates get close to the limit point $X(\infty)$, using techniques akin to those proposed in (Higham, 1999; Fowler & Kelley, 2005; Kelley, Qi, Liao, Reese, & Winton, 2006; Luo, Kelley, Liao, & Tam, 2006), in order to obtain a superlinear algorithm. To this end, it is interesting to investigate how superlinearly convergent methods perform for finding a zero of Oja's vector field (1b). Since Newton's method can be thought of as the prototype superlinear algorithm to which all other superlinear algorithms relate, we propose to investigate the behavior of Newton's method applied to the problem of finding a zero of Oja's vector field (1b).

A crucial hypothesis in the classical local convergence analysis of Newton's method (see, e.g., (Dennis & Schnabel, 1983)) is that the targeted zero is nondegenerate. As it turns out, the zeros of Oja's vector field F (1b) are never isolated, because F displays a property of symmetry under the action of the orthogonal group O_p on the set $\mathbb{R}^{n\times p}$. Therefore, the classical superlinear convergence result of Newton's method is void in the case of F, and indeed our numerical experiments show that Newton's method on $\mathbb{R}^{n\times p}$ for F performs poorly (see Figure 1).

A way to address this difficulty is to modify the vector field F so as to break the symmetry while preserving a useful relation between the zeros of F and the invariant subspaces of A. For example, Oja *et al.* (Oja, Ogawa, & Wangviwattana, 1992b, 1992a) and Xu (Xu, 1993) consider $F : X \mapsto AXD - XDX^TAX$, where D is a nondegenerate positive $p \times p$ diagonal matrix. They prove that p principal eigenvectors (not merely the subspace spanned by them) are derived, and

hence no indefiniteness occurs provided that there is no degeneracy in the eigenvalues of A (i.e., A has n distinct eigenvalues). A detailed analysis appears in (Yoshizawa, Helmke, & Starkov, 2001; Yoshizawa & Helmke, 2002). The algorithm is further extended to derive the minor subspace in a similar form by Chen and Amari (Chen & Amari, 2001), where stability of the algorithm is studied in the geometrical framework.

In this paper, rather than modifying F (1b), we propose a remedy to the degeneracy that consists in "quotienting out" the symmetry of F. Conceptually speaking, instead of performing a Newton iteration in $\mathbb{R}^{n\times p}$, we perform a Newton iteration on a quotient space, namely $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$ (defined in Section 3). We exploit the Riemannian quotient manifold structure of $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$ to formulate a Newton method on this set, following the framework developed in (Absil, Mahony, & Sepulchre, 2008). The resulting Newton equation is a linear matrix equation that can be solved by various numerical approaches. It follows from the theory of Newton method on manifolds (see, e.g., (Adler, Dedieu, Margulies, Martens, & Shub, 2002; Absil, Mahony, & Sepulchre, 2008)), and from a careful analysis of the zeros of the vector field, that the obtained algorithm converges locally superlinearly to the set of orthonormal bases of invariant subspaces of A. Our numerical experiments show that the method behaves as expected.

2 Plain Newton method for Oja's vector field

We assume throughout the paper that A is a real symmetric positive-definite $n \times n$ matrix. For simplicity, we also assume that the eigenvalues of A satisfy

$$\lambda_1 > \cdots > \lambda_n, \tag{2}$$

i.e., all the eigenvalues of A are simple. Hence the p-dimensional invariant subspaces of A are isolated.

Newton's method in $\mathbb{R}^{n\times p}$ for Oja's vector field F (1b) consists of iterating the map $X \mapsto X_+$ defined by solving

$$\mathrm{D}F(X)[Z] \equiv AZ - ZX^TAX - XZ^TAX - XX^TAZ = -F(X) \tag{3}$$

$$X_+ = X + Z. \tag{4}$$

The following proposition is a well-known characterization of the zeros of F.

Proposition 2.1 (zeros of F (1b)) *Let $X \in \mathbb{R}^{n\times p}$ be of full rank. Then the following two conditions are equivalent:*

1. $F(X) = 0$, *i.e.,*
$$AX = XX^TAX. \tag{5}$$

2. $\mathrm{col}(X)$ *is an invariant subspace of A and X is orthonormal ($X^TX = I$).*

Proof. $1 \Rightarrow 2$. We have that $AX = X(X^TAX)$, thus $A\,\mathrm{col}(X) \subseteq \mathrm{col}(X)$. (More precisely, since X has full rank and A is positive-definite, it follows that X^TAX is invertible and thus $A\,\mathrm{col}(X) = \mathrm{col}(X)$.) Moreover, multiplying (5) by X^T on the left yields $X^TAX = X^TXX^TAX$, which implies that $X^TX = I$ since X^TAX is invertible.

$2 \Rightarrow 1$. Since $\mathrm{col}(X)$ is an invariant subspace of A, there is a matrix M such that $AX = XM$. Multiply this equation on the left by X^T to obtain that $M = X^TAX$ and thus (5). □

Hence, the set of all full-rank zeros of F is the finite union of the compact sets

$$\mathcal{S}_i := \{X \in \mathbb{R}^{n\times p} : \mathrm{col}(X) = \mathcal{E}_i, X^TX = I\} \tag{6}$$

where $\mathcal{E}_1, \ldots, \mathcal{E}_N$ are the p-dimensional invariant subspaces of A. (Note that N is finite; it is equal to $\binom{n}{p}$.) It is readily checked that $\mathcal{S}_i = V_i\mathrm{O}_p$, where

$$\mathrm{O}_p := \{Q \in \mathbb{R}^{p\times p} : Q^TQ = I\}$$

is the orthogonal group of order p and V_i is an element of $\mathcal{S}_i$. It follows that the zeros of F are not isolated. Hence the Jacobian of F at the zeros of F is singular (i.e., the zeros are *degenerate*), and consequently, the classical result (see (Dennis & Schnabel, 1983, Th. 5.2.1)) of local superlinear convergence of Newton's iteration to the nondegenerate zeros of F is void. This does not imply that Newton's method will fail, but there is a suspicion that it will behave poorly, and indeed, in Section 5 (see Figure 1), we report on numerical experiments showing that this is the case.

3 Newton's method on $\mathbb{R}^{n\times p}_*/\mathrm{O}_p$

The degeneracy of the zeros of F is due to the following fundamental symmetry property:

$$F(XQ) = F(X)Q, \text{ for all } Q \in \mathrm{O}_p. \tag{7}$$

In this section, we propose a *geometric* Newton method that performs well with functions F that satisfy (7). This geometric Newton method evolves on the quotient space $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$, where this symmetry is removed. Then, in Section 4, we will return to the specific case where F is Oja's vector field (1b) and obtain a concrete numerical algorithm.

The general idea for the geometric Newton method is first to define a vector field ξ on the manifold $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$, whose zeros relate to those of F. The vector field ξ is specified in terms of its so-called *horizontal lift* in $\mathbb{R}_*^{n\times p}$. This formulation requires us to introduce some concepts (vertical and horizontal spaces) borrowed from the theory of fiber bundles (Kobayashi & Nomizu, 1963), or more specifically from the theory of Riemannian submersions (O'Neill, 1983). All the differential-geometric concepts used in this section are explained in (Absil, Mahony, & Sepulchre, 2008), or in (Lee, 2003) for what concerns Lie groups.

The following notation will come useful. Let

$$\mathbb{R}_*^{n\times p} = \{X \in \mathbb{R}^{n\times p} : \det(X^T X) \neq 0\} \tag{8}$$

denote the set of all full-rank $n \times p$ matrices. Let

$$\mathrm{sym}(B) = \frac{1}{2}(B + B^T) \tag{9}$$

and

$$\mathrm{skew}(B) = \frac{1}{2}(B - B^T) \tag{10}$$

denote the terms of the decomposition of a square matrix B into a symmetric term and a skew-symmetric term. For $X \in \mathbb{R}_*^{n\times p}$, define

$$P_X^{\mathrm{p}} : \mathbb{R}^{n\times p} \to \mathbb{R}^{n\times p} : Z \mapsto P_X^{\mathrm{p}}(Z) = (I - X(X^T X)^{-1} X^T)Z \tag{11}$$

$$P_X^{\mathrm{s}} : \mathbb{R}^{n\times p} \to \mathbb{R}^{n\times p} : Z \mapsto P_X^{\mathrm{s}}(Z) = X\mathrm{sym}((X^T X)^{-1} X^T Z) \tag{12}$$

$$P_X^{\mathrm{a}} : \mathbb{R}^{n\times p} \to \mathbb{R}^{n\times p} : Z \mapsto P_X^{\mathrm{a}}(Z) = X\mathrm{skew}((X^T X)^{-1} X^T Z) \tag{13}$$

$$P_X^{\mathrm{h}} = P_X^{\mathrm{p}} + P_X^{\mathrm{s}} : Z \mapsto Z - X\mathrm{skew}((X^T X)^{-1} X^T Z). \tag{14}$$

We have $Z = P_X^{\mathrm{p}}(Z) + P_X^{\mathrm{s}}(Z) + P_X^{\mathrm{a}}(Z)$ for all $Z \in \mathbb{R}^{n\times p}$. Observe that $\mathrm{im}(P_X^{\mathrm{p}}) = \{Z \in \mathbb{R}^{n\times p} : X^T Z = 0\}$, $\mathrm{im}(P_X^{\mathrm{s}}) = X\mathcal{S}_{\mathrm{sym}}(p)$, $\mathrm{im}(P_X^{\mathrm{a}}) = X\mathcal{S}_{\mathrm{skew}}(p)$, where

$$\mathcal{S}_{\mathrm{sym}}(p) = \{S \in \mathbb{R}^{p\times p} : S^T = S\}$$

denotes the set of all symmetric matrices of order p and

$$\mathcal{S}_{\mathrm{skew}}(p) = \{\Omega \in \mathbb{R}^{p\times p} : \Omega^T = -\Omega\}$$

is the set of all *skew-symmetric* (or *antisymmetric*) matrices of order p. (The letter "p" stands for "perpendicular", "s" for symmetric, "a" for antisymmetric, and the notation P^{h} will make sense in a moment.)

In $\mathbb{R}_*^{n\times p}$, we define an equivalence relation "$\sim$" where $X \sim Y$ if and only if there exists a $Q \in \mathrm{O}_p$ such that $Y = XQ$. The equivalence class of $X \in \mathbb{R}_*^{n\times p}$ is thus

$$[X] = X\,\mathrm{O}_p = \{XQ : Q \in \mathrm{O}_p\}. \tag{15}$$

We let

$$\mathbb{R}_*^{n\times p}/\mathrm{O}_p = \mathbb{R}_*^{n\times p}/\sim$$

denote the quotient of $\mathbb{R}_*^{n\times p}$ by this equivalence relation, i.e., the elements of the set $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$ are the equivalence classes of the form (15), $X \in \mathbb{R}_*^{n\times p}$. We let

$$\pi : \mathbb{R}_*^{n\times p} \to \mathbb{R}_*^{n\times p}/\mathrm{O}_p \tag{16}$$

denote the *quotient map* that sends $X \in \mathbb{R}_*^{n\times p}$ to its equivalence class $[X]$ viewed as an element of $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$. The set $\mathbb{R}_*^{n\times p}$ is termed the *total space* of the quotient $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$. Note that a point $\pi(X) \in \mathbb{R}_*^{n\times p}/\mathrm{O}_p$ can be numerically represented by any element of its equivalence class $[X] = \pi^{-1}(\pi(X))$.

Since $\mathbb{R}_*^{n\times p}$ is an open subset of $\mathbb{R}^{n\times p}$, it is naturally an open submanifold of the linear manifold $\mathbb{R}^{n\times p}$. Moreover, it can be shown that the map

$$\psi : \mathrm{O}_n \times \mathbb{R}_*^{n\times p} \to \mathbb{R}_*^{n\times p} : (Q, X) \mapsto XQ$$

is a free and proper Lie group action on the manifold $\mathbb{R}_*^{n\times p}$. Therefore, by the quotient manifold theorem (see, e.g., (Lee, 2003, Th. 9.16)), the orbit space $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$ is a quotient manifold. In other words, the set $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$ is turned into a manifold by endowing it with the unique differentiable structure that makes the quotient map π a submersion. It comes as a consequence that each equivalence class $[X] = \pi^{-1}(\pi(X))$, $X \in \mathbb{R}_*^{n\times p}$, is an embedded submanifold of $\mathbb{R}_*^{n\times p}$. We term *vertical space* at $X \in \mathbb{R}_*^{n\times p}$ the tangent space $\mathcal{V}_X$ to $[X]$ at X, i.e.,

$$\mathcal{V}_X = T_X[X] = X\; T_I\mathrm{O}_p = X\mathcal{S}_{\mathrm{skew}}(p).$$

Observe that $\mathrm{im}(P_X^{\mathrm{a}}) = \mathcal{V}_X$.

Next we define horizontal spaces $\mathcal{H}_X$, which must satisfy the condition that $\mathbb{R}^{n\times p}$ is the internal direct sum of the vertical and horizontal spaces. We choose

$$\mathcal{H}_X = \mathrm{im}(P^{\mathrm{h}}_X) = \{XS + X_\perp K : S^T = S\} \tag{17}$$

where P^{h} is as in (14) and $X_\perp \in \mathbb{R}^{n\times(n-p)}$ denotes any orthonormal matrix such that $X^T X_\perp = 0$. (It would be sufficient to state that $X_\perp$ has full rank, but it does not hurt to assume that it is orthonormal. We do not use $X_\perp$ in the numerical algorithms.)

As an aside, we point out that the horizontal space (17) is the orthogonal complement of the vertical space with respect to the noncanonical metric $\overline{g}$ defined by

$$\overline{g}_X(Z_1, Z_2) = \mathrm{tr}(Z_1^T P^{\mathrm{p}}_X(Z_2) + Z_1^T X(X^T X)^{-2} X^T Z_2).$$

Indeed, we have $\{Z \in \mathbb{R}^{n\times p} : \overline{g}_X(W, Z) = 0 \text{ for all } W \in \mathcal{V}_X\} = \{XS + X_\perp K : S^T = S\}$. The reason for not choosing the canonical Riemannian metric on $\mathbb{R}^{n\times p}_*$ is that it yields a more intricate formula for the horizontal space and for the projection onto it.

Consider a function $F : \mathbb{R}^{n\times p} \to \mathbb{R}^{n\times p}$ that satisfies the symmetry property (7). Define a horizontal vector field $\overline{\xi}$ by

$$\overline{\xi}_X := P^{\mathrm{h}}_X(F(X)).$$

It can be checked that this horizontal vector field $\overline{\xi}$ satisfies $\mathrm{D}\pi(X)[\overline{\xi}_X] = \mathrm{D}\pi(XQ)[\overline{\xi}_{XQ}]$ for all $Q \in \mathrm{O}_p$, and thus $\overline{\xi}$ is the horizontal lift of a vector field ξ on the quotient manifold $\mathbb{R}^{n\times p}_*/\mathrm{O}_p$. In the remainder of this section, we formulate a geometric Newton method for finding a zero of the vector field ξ.

For finding a zero of a vector field ξ on an abstract manifold $\mathcal{M}$ endowed with an *affine connection* ∇ and a *retraction* R, we consider the geometric Newton method in the form proposed by Shub (Adler et al., 2002, §5) (or see (Absil, Mahony, & Sepulchre, 2008, §6.1)). The method consists of iterating the mapping that sends $x \in \mathcal{M}$ to $x_+ \in \mathcal{M}$ obtained by solving

$$\mathrm{J}_\xi(x)[\eta_x] = -\xi_x, \quad \eta_x \in T_x\mathcal{M}, \tag{18a}$$

$$x_+ = R_x(\eta_x). \tag{18b}$$

Here $\mathrm{J}_\xi(x)$ denotes the Riemannian Jacobian $\mathrm{J}_\xi(x) : T_x\mathcal{M} \to T_x\mathcal{M} : \zeta_x \mapsto \mathrm{J}_\xi(x)[\zeta_x] = \nabla_{\zeta_x}\xi$.

For the case where $\mathcal{M}$ is the quotient $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$, we choose the affine connection ∇ defined by

$$\overline{(\nabla_{\eta_{\pi(X)}}\xi)}_X := P_X^{\mathrm{h}}\mathrm{D}\overline{\xi}(X)[\overline{\eta}_X], \tag{19}$$

where the overline denotes the horizontal lift. It can be shown that the right-hand side is indeed a horizontal lift and that the definition of ∇ satisfies all the properties of an affine connection. With this choice for ∇, and with a simple choice for the retraction R (see (Absil, Mahony, & Sepulchre, 2008, §4.1.2) for the relevant theory), the geometric Newton method on the quotient manifold $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$ becomes the iteration that maps $\pi(X)$ to $\pi(X_+)$ by solving

$$P_X^{\mathrm{h}}\mathrm{D}\overline{\xi}(X)[\overline{\eta}_X] = -\overline{\xi}_X, \quad \overline{\eta}_X \in \mathcal{H}_X \tag{20a}$$

$$X_+ = X + \overline{\eta}_X, \tag{20b}$$

with P^{h} as in (14) and $\mathcal{H}_X$ as in (17). Note that, in spite of possibly unfamiliar notation, (20) only involves basic calculus of functions between matrix spaces.

The local superlinear convergence result for the geometric Newton method (20) follows directly from the local convergence result of the general geometric Newton method (18), see (Absil, Mahony, & Sepulchre, 2008, §6.3). It states that the iteration converges locally quadratically to the nondegenerate zeros of ξ. Since we have "quotiented out" the symmetry of F by the action of O_p to obtain ξ, it is now reasonable to hope that the zeros of ξ are nondegenerate. (As we show in the next section, this is indeed the case for Oja's vector field (1b) under the assumption (2) of non-repeated eigenvalues.)

Note that the proposed geometric Newton method differs from the "canonical" Riemannian Newton algorithm (Smith, 1994) on the quotient $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$, because the affine connection chosen is not the Riemannian connection on $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$ endowed with the Riemannian metric inherited from the canonical metric in $\mathbb{R}_*^{n\times p}$. However, the property of local superlinear convergence to the nondegenerate zeros still holds (see (Adler et al., 2002) or (Absil, Mahony, & Sepulchre, 2008)).

4 A geometric Newton method for Oja's vector field

In this section, we apply the geometric Newton method on $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$, given in (20), to the case where the tangent vector field ξ on $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$ is defined by the horizontal lift

$$\overline{\xi}_X := P_X^{\mathrm{h}}(F(X)), \tag{21}$$

with P^{h} as in (14) and F as in (1b). The resulting Newton iteration is formulated in Algorithm 1. Recall the definitions of $\mathbb{R}_*^{n\times p}$ (8), P^{h} (14), skew (10), $\mathcal{H}_X$ (17).

Algorithm 1 Geometric Newton for Oja's vector field

Require: Symmetric positive-definite $n \times n$ matrix A; positive integer $p < n$.

Input: Initial iterate $X_0 \in \mathbb{R}_*^{n\times p}$, i.e., X_0 is a real $n \times p$ matrix with full rank.

Output: Sequence of iterates $(X_k) \subset \mathbb{R}_*^{n\times p}$.

1: **for** $k = 0, 1, 2, \ldots$ **do**

2: Solve the linear system of equations (we drop the subscript k for convenience)

$$\begin{aligned} &P_X^{\mathrm{h}}(AZ - ZX^TAX - XZ^TAX - XX^TAZ - Z\mathrm{skew}((X^TX)^{-1}X^TAX) \\ &- X\mathrm{skew}(-(X^TX)^{-1}(X^TZ + Z^TX)(X^TX)^{-1}X^TAX + (X^TX)^{-1}(Z^TAX + X^TAZ))) \\ &\qquad = -\big(AX - XX^TAX - X\mathrm{skew}((X^TX)^{-1}X^TAX)\big) \end{aligned} \tag{22}$$

for the unknown $Z \in \mathcal{H}_X$.

3: Set

$$X_{k+1} = X_k + Z.$$

4: **end for**

Observe that Algorithm 1 is stated as an iteration in the total space $\mathbb{R}_*^{n\times p}$ of the quotient $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$. Formally, the sequence of iterates of the Newton method on $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$, for an initial point $\pi(X_0) \in \mathbb{R}_*^{n\times p}/\mathrm{O}_p$, is given by $(\pi(X_k))_{k\in\mathbb{N}}$, where π is the quotient map (16) and $(X_k)_{k\in\mathbb{N}}$ is the sequence of iterates generated by Algorithm 1.

We point out that (22) is merely a linear system of equations. It can be solved by means of iterative solvers that can handle linear systems in operator form. Moreover, these solvers can be stopped early to avoid unnecessary computational effort when the iterate X_k is still far away from the solution. Guidelines for stopping the linear system solver can be found,

e.g., in (Eisenstat & Walker, 1996). In our numerical experiments (Section 5) we have used MATLAB's GMRES solver.

Algorithm 1 converges locally quadratically to the nondegenerate zeros of ξ. We first characterize the zeros of ξ, then we show that they are all nondegenerate under the assumption (2).

First note that $\xi_{\pi(X)} = 0$ if and only if $P^{\mathrm{h}}_X(F(X)) = 0$, where $X \in \mathbb{R}^{n\times p}_*$. Under the assumption that $X \in \mathbb{R}^{n\times p}_*$, the following statements are equivalent:

$$P^{\mathrm{h}}_X(F(X)) = 0,$$

$$F(X) \in \mathrm{im}(P^{\mathrm{a}}_X),$$

$$F(X) = X\Omega \text{ for some } \Omega = -\Omega^T,$$

$$AX - XX^TAX = X\Omega \text{ for some } \Omega = -\Omega^T,$$

$$A\,\mathrm{col}(X) \subseteq \mathrm{col}(X) \text{ and } (X^TX)^{-1}X^TAX - X^TAX \text{ is skew-symmetric},$$

$$A\,\mathrm{col}(X) \subseteq \mathrm{col}(X) \text{ and } (X^TX)^{-1}X^TAX + X^TAX(X^TX)^{-1} = 2X^TAX.$$

We thus have an equation of the form $YB + BY = 2B$, where $B := X^TAX$ is symmetric positive definite, hence its eigenvalues β_i, $i = 1, \ldots, p$, are all real and positive. The Sylvester operator $Y \mapsto YB + BY$ is a linear operator whose eigenvalues are $\beta_i + \beta_j$, $i = 1, \ldots, p$, $j = 1, \ldots, p$ (Gantmacher, 1959, Ch. VI). All these eigenvalues are real and positive, thus nonzero, hence the operator is nonsingular, from which it follows that the equation $YB + BY = 2B$ has one and only one solution Y. It is readily checked that this unique solution is $Y = I$. We have thus shown that $X^TX = I$. The result is summarized in the following proposition.

Proposition 4.1 *Let P^{h} be as in* (14) *and let F be Oja's vector field* (1b). *Then $X \in \mathbb{R}^{n\times p}_*$ is a zero of the projected Oja vector field $P^{\mathrm{h}}(F)$ if and only if $A\,\mathrm{col}(X) \subseteq \mathrm{col}(X)$ and $X^TX = I$.*

In other words, the full-rank zeros of the projected Oja vector field $P^{\mathrm{h}}(F)$ are the full-rank zeros of Oja's vector field F. This means that we do not lose information by choosing to search for the zeros of ξ (21) instead of F (1b).

It remains to show that the zeros of ξ are nondegenerate. Let $\pi(X_*)$ be a zero of ξ, which means that $AX_* = X_*X_*^TAX_*$ and $X_*^TX_* = I$. The task is to show that the Jacobian operator $\mathrm{J}_\xi(\pi(X_*))$, or equivalently its lifted counterpart

$$\overline{\mathrm{J}_\xi}(X_*) : \mathrm{im}(P^{\mathrm{h}}) \to \mathrm{im}(P^{\mathrm{h}}) : Z \mapsto P^{\mathrm{h}}_{X_*}\left(\mathrm{D}(P^{\mathrm{h}}(F))(X_*)[Z]\right), \tag{23}$$

is nonsingular. To this end, consider the operator

$$\mathcal{J} : \mathbb{R}^{n\times p} \to \mathbb{R}^{n\times p} : Z \mapsto \mathrm{D}(P^{\mathrm{h}}(F))(X_*)[Z]. \tag{24}$$

Note that $\overline{\mathrm{J}_\xi}(X_*)$ is the restriction of $\mathcal{J}$ to $\mathrm{im}(P^{\mathrm{h}})$. Consider the decomposition $\mathbb{R}^{n\times p} = \mathrm{im}(P^{\mathrm{p}}) \oplus \mathrm{im}(P^{\mathrm{s}}) \oplus \mathrm{im}(P^{\mathrm{a}})$ and recall that $\mathrm{im}(P^{\mathrm{h}}) = \mathrm{im}(P^{\mathrm{p}}) \oplus \mathrm{im}(P^{\mathrm{s}})$. We show that the corresponding block decomposition of $\mathcal{J}$ is as follows:

$$\begin{bmatrix} * & 0 & 0 \\ ? & * & 0 \\ ? & ? & 0 \end{bmatrix},$$

where "*" denotes nonsingular operators. It then directly follows that the upper two-by-two block of $\mathcal{J}$, which corresponds to $\overline{\mathrm{J}_\xi}(X_*)$, is nonsingular.

We show that the 11 block (i.e., the "pp" block) is nonsingular. This block is the operator from $\mathrm{im}(P^{\mathrm{p}}_{X_*})$ to $\mathrm{im}(P^{\mathrm{p}}_{X_*})$ given by

$$Z \;\mapsto\; P^{\mathrm{p}}_{X_*}\mathrm{D}(P^{\mathrm{h}}(F))(X_*)[Z] = P^{\mathrm{p}}_{X_*}AZ - P^{\mathrm{p}}_{X_*}ZX_*^TAX_*.$$

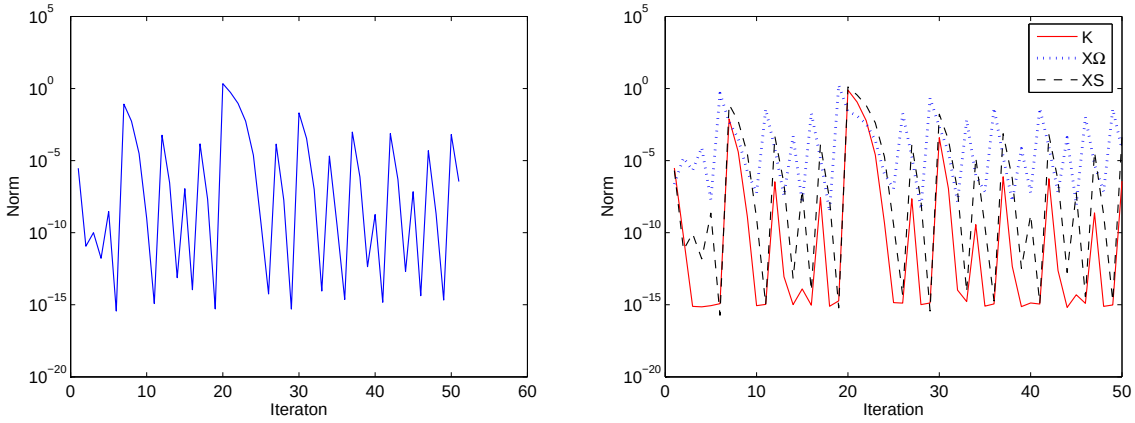
(In obtaining this result, we have used the relations $X_*^TX_* = I$, $\mathrm{skew}(X_*^TAX_*) = 0$, $P^{\mathrm{p}}_{X_*}X_* = 0$.) In view of the hypothesis that the eigenvalues of A are all simple, this operator is known to be nonsingular; see, e.g., (Lundström & Eldén, 2001/02; Absil, Sepulchre, Van Dooren, & Mahony, 2004).

We show that the 22 block (i.e., the "ss" block) is nonsingular. This block is the operator from $\mathrm{im}(P^{\mathrm{s}}_{X_*})$ to $\mathrm{im}(P^{\mathrm{s}}_{X_*})$

$$X_*S \quad\mapsto\quad P^{\mathrm{s}}_{X_*}\Big(\mathrm{D}(P^{\mathrm{h}}(F))(X_*)[X_*S]\Big) = -X_*(SX_*^TAX_* + X_*^TAX_*S).$$

(We have used the relation $AX_* = X_*X_*^TAX_*$ to obtain this expression.) In view of the previous discussion on the Sylvester operator, this operator is nonsingular.

This completes the proof that the zeros of ξ are nondegenerate. Consequently, for all X_0 sufficiently close to some $\mathcal{S}_i$ (6), the sequence (X_k) generated by Algorithm 1 is such that $X_k\mathrm{O}_p$ converges quadratically to $\mathcal{S}_i$. Recall that $\mathcal{S}_i$ is the set of all orthonormal matrices whose column space is the ith invariant subspace of A.



(a) F

(b) K, XS, $X\Omega$

Figure 1: Plain Newton method

5 Numerical experiments

In this section, we report on numerical experiments for both the plain Newton and the geometric Newton method, derived in Section 2 and Section 4, respectively. The experiments were run using MATLAB. The machine epsilon is approximately $2 \cdot 10^{-16}$.

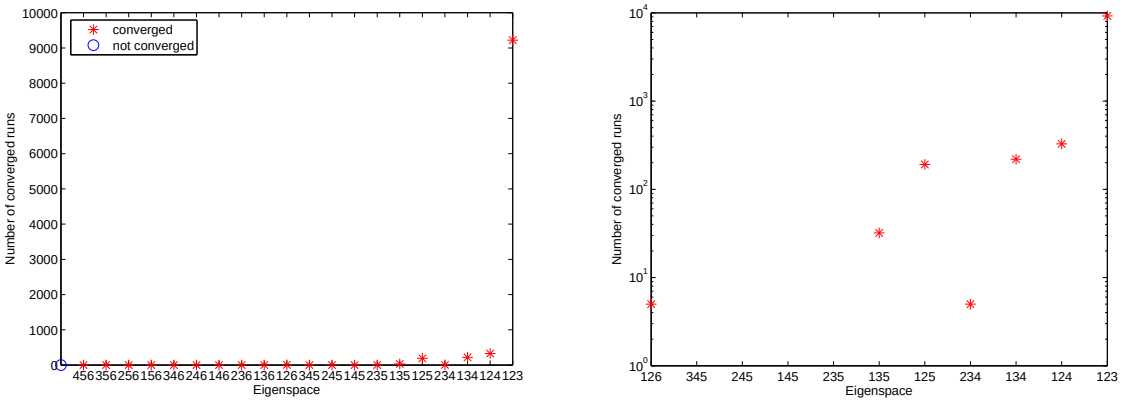
As mentioned in Section 1, the plain Newton method performs poorly due to the fact that the zeros of the cost function F are not isolated. To illustrate this, we consider a symmetric positive definite matrix $A \in \mathbb{R}^{6\times 6}$ with uniformly distributed eigenvalues in the interval $[0, 1]$ and 100 different initial iterates $X_0 \in \mathbb{R}^{6\times 3}$. Each of the initial iterates is computed as

$$X_0 = X_* + 10^{-6} E \,,$$

where X_* is such that $F(X_*) = 0$ and E is a 6×3 matrix with random entries, chosen from a normal distribution with zero mean and unit variance. The simulation is stopped after 50 iterations. One representative example is given in Figure 1. In Figure 1 (a), the norm of $F(X)$ is given for each iteration step. Close to the solution, the system matrix gets singular and the algorithm deviates from the optimal point. In Figure 1 (b), we present the evolution of the norms of the three components $K, X\Omega$, and XS of the update vector Z,

$$Z = X_\perp K + X\Omega + XS \,,$$

where $X_\perp$ is the orthogonal complement of X, Ω is a skew-symmetric matrix and S is a symmetric matrix. We see that, even when the K and XS component are very small, $X\Omega$ is



(a) Linear plot

(b) Logarithmic scale

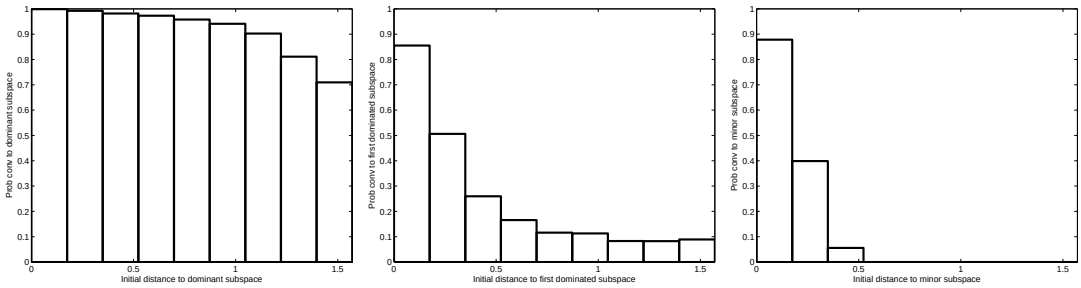
Figure 2: Geometric Newton method

quite large. This concords with the fact that the kernel of the Hessian at a stationary point X is $\{X\Omega : \Omega^T = -\Omega\}$.

Next, we study the geometric Newton method, derived in Section 4. Again, we consider $n = 6$, $p = 3$ and a symmetric positive definite matrix $A \in \mathbb{R}^{6\times 6}$ with uniformly distributed eigenvalues in the interval $[0, 1]$. We perform 10^4 experiments with a single matrix A but with different initial iterates $X_0 \in \mathbb{R}^{6\times 3}_*$ with random entries, chosen from a normal distribution with zero mean and unit variance. In Figure 2, we show the number of runs that converged to each of the eigenspaces of A. The dominant eigenspace is marked by "123". In general, "ijk" stands for the eigenspace spanned by the i^{th}, j^{th}, and k^{th} eigenvectors [1]. Let W be the matrix of all eigenvectors. We declare that the algorithm has converged to eigenspace "ijk" when the largest principal angle (Golub & Van Loan, 1996, §12.4.3) between this eigenspace and X is smaller than 10^{-10} after 50 iterations. In our experiment, all the runs have converged to one of the 20 possible eigenspaces. In general, there may be cases where the algorithm does not converge to any of the eigenspaces. This may occur when the initial iterate X_0 is very close to the boundary of one of the basins of attraction. However, these cases are rare.

It appears that the basin of attraction of the dominant eigenspace is the largest. This finding was not expected *a priori*, but it is not in contradiction with any known property of Algorithm 1. Note in particular that, in contrast with the Grassmann-based Newton method for the

[1] For convenience, we consider an eigenvalue decomposition, where the eigenvalues of A are put in descending order.



(a) Dominant eigenspace (b) First dominated eigenspace (c) Minor eigenspace

Figure 3: Distribution of the basin of attraction of three eigenspaces of interest. The horizontal axis shows the largest principal angle between $\mathrm{col}(X(0))$ and the eigenspace under consideration. The vertical axis gives the fraction of initial conditions that led to convergence to the eigenspace.

generalized Rayleigh quotient investigated in (Edelman, Arias, & Smith, 1998; Lundström & Eldén, 2001/02; Absil, Mahony, & Sepulchre, 2004), Algorithm 1 is not invariant by reflections of the spectrum of A; namely, Algorithm 1 is not invariant by transformations of the type $A \mapsto -(A-\mu I)$, $\mu \in \mathbb{R}$. Therefore, whereas the Grassmann-based method does not favor one end of the spectrum over the other, there is no reason why Algorithm 1 should display a similar behavior. Depending on the application, the observed global behavior of Algorithm 1, with a large basin of attraction for the dominant eigenspace, can be considered an advantage over the Grassmann-based iteration.

To investigate this issue further, we performed numerical experiments to assess the distribution of the area of the various basins of attraction, measured in terms of the largest principal angle between $\mathrm{col}(X(0))$ and eigenspaces of interest. It is known (see, e.g., (Absil, Edelman, & Koev, 2006)) that if the elements of $X(0)$ are drawn from a Gaussian distribution, then the bins corresponding to small values of the largest principal angle are scarcely populated, hence the statistics for these bins are unreliable. (To remedy this, it is possible to draw $X(0)$ from a Gaussian distribution and discard some of them according to the probability density function given in (Absil et al., 2006), so as to obtain a uniform distribution on a prescribed subinterval of $[0, \pi/2]$ for the value of the largest principal angle. However, to obtain a uniform distribution on a reasonably large interval, we had to discard so many $X(0)$'s that this approach turned out to be impractical.) Therefore, in the numerical experiments mentioned below, we drew $X(0)$ according to $X(0) = X_* + aE$, where a is uniformly distributed on the unit interval and the elements of E are drawn from the Gaussian distribution with mean zero and standard deviation one; this

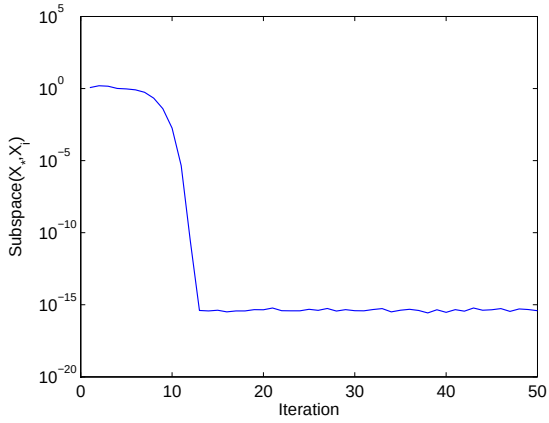


Figure 4: Superlinear convergence of the geometric Newton method

technique resulted in bins that are reasonably evenly populated. We declare that the iteration has converged to eigenspace $\mathcal{E}_i$ if the largest principal angle between $\mathrm{col}(X(k))$ and $\mathcal{E}_i$ is smaller than 10^{-10} when $k = 50$. The distribution of the area of the basins of attraction corresponding to the dominant eigenspace, to the first dominated eigenspace (i.e., eigenspace 124), and to the minor eigenspace, are shown on the histograms of Figure 3, obtained over 10^4 experiments. We observed that the sequences that do not converge to the dominant eigenspace tend to display a rather large first step, which suggests that they are initialized close to a point where the Jacobian is not invertible.

In interpreting the results of Figure 3, one should bear in mind that Algorithm 1 induces an iteration on $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$, but not on the Grassmann manifold. Hence, the largest principal angle to $\mathrm{col}(X(0))$ is not a measure of distance to the solution; for example, if the column space of $X(0)$ is close to the dominant eigenspace but $X(0)$ is far from orthonormality, then $X(0)$ is far away from any zero of Oja's vector field even though $X(0)$ may be tallied in the leftmost bin of the first histogram.

Finally, the superlinear convergence rate of the algorithm is illustrated in Figure 4.

6 Conclusion

We have investigated the use of Newton's method to compute superlinearly the zeros of Oja's vector field (1b). Due to a symmetry in the vector field by the action of the orthogonal group,

its zeros are never isolated, which causes the plain Newton method to behave poorly. We have proposed a remedy that consists in "quotienting out" the symmetry. This led to the formulation of a geometric Newton algorithm that seeks the zeros of a projection of Oja's vector field. We have shown that the zeros of the projected vector field are the same as the zeros of the original vector field. Moreover, these zeros are nondegenerate. This means that by quotienting out the action of the orthogonal group, we have removed just enough symmetry to make the zeros nondegenerate. In view of the nondegeneracy property, it follows directly from the convergence theory of the abstract geometric Newton method that the resulting algorithm converges locally superlinearly to the zeros of Oja's vector field.

Invariant subspace computation has been and still is a very active area of research. As a method for invariant subspace computation, it is doubtful that the proposed algorithm can outperform the state-of-the-art methods. In particular, the Grassmann-based approach (Edelman et al., 1998; Lundström & Eldén, 2001/02; Absil, Mahony, & Sepulchre, 2004), that can be thought of as quotienting out the action of the whole general linear group instead of the orthogonal group, leads to a Newton equation that lives in a smaller subspace of $\mathbb{R}^{n\times p}$ and that can be solved in fewer flops. When $n \gg p$, however, the number of operations necessary to apply the linear operator in the Newton equation is of the same order. Moreover, the Grassmann-based algorithm does not possess the remarkable feature of naturally converging towards orthonormal matrices, i.e., without having to resort to orthogonalization steps such as Gram-Schmidt.

The problem of computing the zeros of Oja's vector field was also an occasion for introducing the quotient manifold $\mathbb{R}_*^{n\times p}/\mathrm{O}_p$, that seems to have received little attention in the literature, in contrast to the more famous Grassmann and Stiefel manifolds. In later work, we will further analyze the geometry of this quotient manifold, which was just touched upon in this paper, and we will show how it can be used in the context of low-rank approximation problems.

Acknowledgements

Fruitful discussions with Michel Journée, Rodolphe Sepulchre, and Bart Vandereycken are gratefully acknowledged.

References

Absil, P.-A., Edelman, A., & Koev, P. (2006). On the largest principal angle between random subspaces. *Linear Algebra Appl.*, *414*(1), 288–294.

Absil, P.-A., Mahony, R., & Sepulchre, R. (2004, January). Riemannian geometry of Grassmann manifolds with a view on algorithmic computation. *Acta Appl. Math.*, *80*(2), 199–220.

Absil, P.-A., Mahony, R., & Sepulchre, R. (2008). *Optimization algorithms on matrix manifolds*. Princeton, NJ: Princeton University Press.

Absil, P.-A., Sepulchre, R., & Mahony, R. (2008). Continuous-time subspace flows related to the symmetric eigenproblem. *Pac. J. Optim.*, *4*(2), 179–194.

Absil, P.-A., Sepulchre, R., Van Dooren, P., & Mahony, R. (2004). Cubically convergent iterations for invariant subspace computation. *SIAM J. Matrix Anal. Appl.*, *26*(1), 70–96.

Adler, R. L., Dedieu, J.-P., Margulies, J. Y., Martens, M., & Shub, M. (2002, July). Newton's method on Riemannian manifolds and a geometric model for the human spine. *IMA J. Numer. Anal.*, *22*(3), 359–390.

Chen, T. P., & Amari, S. (2001). Unified stabilization approach to principal and minor components extraction algorithms. *Neural Networks*, *14*(10), 1377–1387.

Dehaene, J., Moonen, M., & Vandewalle, J. (1999). Analysis of a class of continuous-time algorithms for principal component analysis and subspace tracking. *IEEE Trans. Circuits Systems I Fund. Theory Appl.*, *46*(3), 364–372.

Dennis, J. E., Jr., & Schnabel, R. B. (1983). *Numerical methods for unconstrained optimization and nonlinear equations*. Englewood Cliffs, NJ: Prentice Hall Inc.

Edelman, A., Arias, T. A., & Smith, S. T. (1998). The geometry of algorithms with orthogonality constraints. *SIAM J. Matrix Anal. Appl.*, *20*(2), 303–353.

Eisenstat, S., & Walker, H. (1996). Choosing the forcing terms in an inexact Newton method. *SIAM J. Sci. Comput.*, *17*(1), 16-32.

Fowler, K. R., & Kelley, C. T. (2005). Pseudo-transient continuation for nonsmooth nonlinear equations. *SIAM J. Numer. Anal.*, *43*(4), 1385–1406 (electronic).

Gantmacher, F. R. (1959). *The theory of matrices I, II*. New-York: Chelsea.

Golub, G. H., & Van Loan, C. F. (1996). *Matrix computations* (Third ed.). Baltimore, MD:

Johns Hopkins University Press.

Higham, D. J. (1999). Trust region algorithms and timestep selection. *SIAM J. Numer. Anal.*, *37*(1), 194–210 (electronic).

Jankovic, M. V., & Ogawa, H. (2006). Modulated Hebb–Oja learning rule—a method for principal subspace analysis. *IEEE Trans. Neural Networks*, *17*(2), 345–356.

Kelley, C. T., Qi, L., Liao, L.-Z., Reese, J. P., & Winton, C. (2006). *Pseudo-transient continuation and optimization.* (http://www4.ncsu.edu/∼ctk/pubs.html)

Kobayashi, S., & Nomizu, K. (1963). *Foundations of differential geometry.* Interscience Publishers, a division of John Wiley & Sons, New York-London. (Volumes 1 and 2)

Lee, J. M. (2003). *Introduction to smooth manifolds* (Vol. 218). New York: Springer-Verlag.

Lundström, E., & Eldén, L. (2001/02). Adaptive eigenvalue computations using Newton's method on the Grassmann manifold. *SIAM J. Matrix Anal. Appl.*, *23*(3), 819–839.

Luo, X.-L., Kelley, C. T., Liao, L.-Z., & Tam, H.-W. (2006). *Combining trust region techniques and Rosenbrock methods for gradient systems.* (http://www4.ncsu.edu/∼ctk/pubs.html)

Manton, J. H., Helmke, U., & Mareels, I. M. Y. (2005). A dual purpose principal and minor component flow. *Systems Control Lett.*, *54*(8), 759–769.

Oja, E. (1982). A simplified neuron model as a principal component analyzer. *J. Math. Biol.*, *15*, 267-273.

Oja, E. (1989). Neural networks, principal components, and subspaces. *Int. J. Neural Syst.*, *1*, 61–68.

Oja, E., Ogawa, H., & Wangviwattana, J. (1992a). Principal component analysis by homogeneous neural networks, Part II: Analysis and extensions of the learning algorithms. *IEICE Trans. Inf. Syst.*, *3*, 376-382.

Oja, E., Ogawa, H., & Wangviwattana, J. (1992b). Principal component analysis by homogeneous neural networks, Part I: The weighted subspace criterion. *IEICE Trans. Inf. Syst.*, *3*, 366-375.

O'Neill, B. (1983). *Semi-Riemannian geometry* (Vol. 103). New York: Academic Press Inc. [Harcourt Brace Jovanovich Publishers].

Ortega, J. M., & Rheinboldt, W. C. (1970). *Iterative solution of nonlinear equations in several variables.* New York: Academic Press.

Smith, S. T. (1994). Optimization techniques on Riemannian manifolds. In *Hamiltonian and gradient flows, algorithms and control* (Vol. 3, pp. 113–136). Providence, RI: Amer. Math. Soc.

Xu, L. (1993). Least mean square error reconstruction principle for self-organizing neural-nets. *Neural Networks*, *6*(5), 627–648.

Yan, W.-Y. (1998). Stability and convergence of principal component learning algorithms. *SIAM J. Matrix Anal. Appl.*, *19*(4), 933–955 (electronic).

Yan, W.-Y., Helmke, U., & Moore, J. B. (1994). Global analysis of Oja's flow for neural networks. *IEEE Trans. on Neural Networks*, *5*(5), 674-683.

Yoshizawa, S., & Helmke, U. (2002). *Corrections to convergence analysis for principal component flows.* (preprint)

Yoshizawa, S., Helmke, U., & Starkov, K. (2001). Convergence analysis for principal component flows. *Int. J. Appl. Math. Comput. Sci.*, *11*(1), 223-236.