

Minimum-Correction Second-Moment Matching:
Theory, Algorithms and Applications

by

Jing Lin

B.Eng., Tsinghua University (2012)

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Master of Science in Mechanical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2020

© Massachusetts Institute of Technology 2020. All rights reserved.

Signature redacted


Author
Department of Mechanical Engineering
Sep 12th, 2019

Signature redacted

Certified by
Pierre F.J. Lermusiaux
Professor, Department of Mechanical Engineering
Thesis Supervisor

Signature redacted

Accepted by
Nicolas Hadjiconstantinou
Chairman, Department Committee on Graduate Theses





77 Massachusetts Avenue
Cambridge, MA 02139
<http://libraries.mit.edu/ask>

DISCLAIMER NOTICE

Due to the condition of the original material, there are unavoidable flaws in this reproduction. We have made every effort possible to provide you with the best copy available.

Thank you.

The images contained in this document are of the best quality available.

Minimum-Correction Second-Moment Matching: Theory, Algorithms and Applications

by

Jing Lin

Submitted to the Department of Mechanical Engineering
on Sep 12th, 2019, in partial fulfillment of the
requirements for the degree of
Master of Science in Mechanical Engineering

Abstract

We address the problem of finding the closest matrix \tilde{U} to a given U under the constraint that a prescribed second-moment matrix \tilde{P} must be matched, i.e. $\tilde{U}^T \tilde{U} = \tilde{P}$. We obtain a closed-form formula for the unique global optimizer \tilde{U} for the full-rank case, which is related to U by an SPD (symmetric positive definite) linear transform. This result is generalized to rank-deficient cases as well as to infinite dimensions. We highlight the geometric intuition behind the theory and study the problem's rich connections to minimum congruence transform, generalized polar decomposition, optimal transport, and rank-deficient data assimilation. In the special case of $\tilde{P} = I$, minimum-correction second-moment matching reduces to the well-studied optimal orthonormalization problem. We investigate the general strategies for numerically computing the optimizer, analyze existing polar decomposition and matrix square root algorithms. More importantly, we modify and stabilize two Newton iterations previously deemed unstable for computing the matrix square root, which can now be used to efficiently compute both the orthogonal polar factor and the SPD square root. We then verify the higher performance of the various new algorithms using benchmark cases with randomly generated matrices. Lastly, we complete two applications for the stochastic Lorenz-96 dynamical system in a chaotic regime. In reduced subspace tracking using dynamically orthogonal equations, we maintain the numerical orthonormality and continuity of time-varying base vectors. In ensemble square root filtering for data assimilation, the prior samples are transformed into posterior ones by matching the covariance given by the Kalman update while also minimizing the corrections to the prior samples.

Thesis Supervisor: Pierre F.J. Lermusiaux
Title: Professor, Department of Mechanical Engineering

Acknowledgments

This thesis would not have been possible without the help and inspiration from many people. First and foremost, I am very grateful to my advisor, Professor Pierre Lermusiaux, for pointing me to this problem, encouraging me to overcome the difficulties, providing insightful feedback and funding my study and research over the years. I am grateful to the Office of Naval Research for support under grants N00014-14-1-0476 (Science of Autonomy – LEARNS) and N00014-14-1-0725 (Bays-DA) to the Massachusetts Institute of Technology.

Several members of our MSEAS group also contributed to this work. I want to thank Florian Feppon for several useful suggestions and for pointing me to the literature on the polar decomposition, which led to a leap of the work. I am also grateful to Deepak Subramani and Arkopal Dutt for many inspiring discussions on re-orthonormalization for dynamically orthogonal (DO) equations. I am delighted that an early result of this work helped them clear an obstinate obstacle in applying multi-step time-marching schemes to DO. Thanks also to all the MSEAS group members that I have ever met, for the friendship and help they offer.

I also want to thank Leslie Regan, our academic administrator in MechE, for a lot of help received from her. Without her heart for students and her devotion to work, the whole MechE experience would be much more difficult and much less enjoyable.

Lastly, I would like to express my deepest gratitude to my parents, for their sacrifice in bringing me up and for their unconditional love and care throughout my journey of life.

Contents

1	Introduction	13
2	Theory of Minimum-Correction Second-Moment Matching	17
2.1	Analytic Expression of the Global Optimizer	17
2.2	Minimum Linear Operation for Second-Moment Matching and Minimum Congruence Transform	23
2.3	Minimum-Correction Orthonormalization and Polar Decomposition with Generalization	25
2.4	Generalization from \mathbb{R}^m to a Hilbert Space	26
2.5	Connections to Optimal Transport	30
2.6	Complications due to Degeneracy	34
2.7	Maximization Counterparts	39
3	Algorithms for Minimum-Correction Second-Moment Matching	41
3.1	General Computation Strategies	41
3.2	Algorithms for the SPD Square Root and Polar Decomposition	43
3.2.1	Algorithms Based on SVD/EVD	43
3.2.2	Algorithms Based on Fixed-Point Iterations	44
3.3	The Special Case of Orthonormalization	51
4	Benchmarks of the Algorithms' Performance	55
5	Application to Matrix Differential Equations that Preserve Orthogonality	59
6	Application to Ensemble Square Root Filters for Data Assimilation	65

7	Conclusions	69
A	Computing the Matrix Gradients	71
B	Successive First Order Approximation for Matrix Square Root	73
C	Gradient Descent for Matrix Square Root	75

List of Figures

2-1	A geometric picture of theorem 2.1 for $m = 2$ and $n = 1$	22
5-1	Orthogonality errors in the absence of re-orthonormalization. The error saturates at a higher level with larger time step Δt . When $\Delta t = 0.05$, the orthogonality is completely lost for $n = 10$ (left), while the solution even blows up for $n = 20$ (right).	62
5-2	Comparison between orthonormalization using SVD-EVD (left) with that using polar decomposition PD-EVD (right). The evolution (from $t = 10$ to 50, $\Delta t = 0.05$) of the first two DO modes $[\mathbf{u}_1^T, \mathbf{u}_2^T]^T$ stacked together are shown on the top and the evolution of their coefficient variances $\text{Var}(\phi_1)$ and $\text{Var}(\phi_2)$ on the bottom.	63
5-3	The orthogonality errors after re-orthonormalization by various algorithms (left) and the ratio of the correction $\ \tilde{\mathbf{U}} - \hat{\mathbf{U}}\ $ to the minimum $\ \tilde{\mathbf{U}}_* - \hat{\mathbf{U}}\ $ computed by (2.11) (right). The ratios for SVD-EVD are of order 10^6 and thus not plotted. Here $\Delta t = 0.05$ and $n = 10$ are used for all the simulations.	64
6-1	Time evolution of RMSE of the various filters (left) and the ratio of the filtering correction $\ \tilde{\mathbf{U}}_+ - \tilde{\mathbf{U}}_-\ $ to the minimum $\ \tilde{\mathbf{U}}_{+*} - \tilde{\mathbf{U}}_-\ $ computed by (2.17) (right). The first three cases are without dimension reduction while the last one is reduced to a 10-dimensional subspace using DO. Here $\Delta t = 0.05$ is used for all the simulations.	67

List of Tables

3.1	Recommended algorithms for computing the minimum-correction second-moment matching solution \tilde{U}_* in theorem 2.1. \tilde{U}_* -based algorithms avoid forming $P = U^T U$ and directly operate on U	51
4.1	Description of the various orthonormalization algorithms benchmarked. Here U is to be orthonormalized and $P = U^T U$ is its Gram matrix. The algorithms that do not form P and directly operate on U are marked by a “*” in their names.	56
4.2	Performance of the various algorithms for orthonormalization in the case of well-conditioned (nearly orthogonal) rectangular matrices (m -by- n , $m \gg n$). Here A is the linear transform such that $\tilde{U} = UA$. The Frobenius norm is exclusively used. For the algorithms based on a fixed-point iteration, the average number of iterations is recorded in the parenthesis in front of the run time. In this case, we have $\ U^T U - I\ = 6.5$ for all 10 repetitions, which can be compared to $\ \tilde{U}^T \tilde{U} - I\ $. . .	57
4.3	Performance of the algorithms on well-conditioned square matrices ($\kappa(U) = 1.5$, $\ U^T U - I\ = 29$) and on ill-conditioned rectangular matrices ($\kappa(U) = 10^6$, $\ U^T U - I\ = 1.5 \times 10^{12}$), respectively. Notations are as those in table 4.2.	57

Chapter 1

Introduction

The second-moment matrix of $U \in \mathcal{M}_{m \times n}(\mathbb{R})$ ¹ defined as $P = U^T \Gamma_m U \in \mathcal{M}_{n \times n}$ contains all the pairwise inner products of U 's columns weighted by Γ_m . When we view U 's columns as vectors in an inner product space, P is called the Gram matrix. When the U 's rows are samples of a zero-mean random vector, P becomes the sample covariance matrix with $\Gamma_m = (1/m)\mathbf{I}$. Here we unify them as second-moment matrices. In many computational problems, we encounter the task of correcting a given matrix U to some \tilde{U} that matches a prescribed second-moment matrix \tilde{P} , i.e. $\tilde{P} = \tilde{U}^T \Gamma_m \tilde{U}$. However, such a \tilde{U} is not unique. When there is no physical information that favors one choice over another, a natural approach is to aim for a correction $\tilde{U} - U$ that is minimal to avoid introducing numerical artifacts.

Precisely, given a target symmetric positive definite (SPD) second-moment matrix \tilde{P} , as well as an inner product on \mathbb{R}^m and on \mathbb{R}^n defined by the SPD weight matrix $\Gamma_m \in \mathcal{M}_{m \times m}$ and $\Gamma_n \in \mathcal{M}_{n \times n}$, respectively, we want to solve the optimization

$$\arg \min_{\tilde{U} \in \mathcal{M}_{m \times n}: \tilde{U}^T \Gamma_m \tilde{U} = \tilde{P}} \|\tilde{U} - U\|_{\mathbb{F}, \Gamma_m, \Gamma_n}^2, \quad (1.1)$$

¹Since we will exclusively focus on real matrices in this paper, the explicit specification “ (\mathbb{R}) ” of the field of matrix elements will be dropped hereafter. However, everything discussed here readily generalizes to the complex case.

where $\|\cdot\|_{\mathbf{F}, \mathbf{\Gamma}_m, \mathbf{\Gamma}_n}$ is the Frobenius norm weighted by $\mathbf{\Gamma}_m$ and $\mathbf{\Gamma}_n$:

$$\begin{aligned}\|\mathbf{V}\|_{\mathbf{F}, \mathbf{\Gamma}_m, \mathbf{\Gamma}_n} &\triangleq \text{tr}(\sqrt{\mathbf{\Gamma}_n} \mathbf{V}^T \mathbf{\Gamma}_m \mathbf{V} \sqrt{\mathbf{\Gamma}_n}^T) \\ &= \text{tr}(\sqrt{\mathbf{\Gamma}_m} \mathbf{V} \mathbf{\Gamma}_n \mathbf{V}^T \sqrt{\mathbf{\Gamma}_m}^T) = \|\mathbf{V}^T\|_{\mathbf{F}, \mathbf{\Gamma}_n, \mathbf{\Gamma}_m}.\end{aligned}\tag{1.2}$$

Here $\sqrt{\mathbf{A}}$ denotes a matrix square root of an SPD matrix \mathbf{A} , viewed as a self-adjoint positive operator, i.e. $\sqrt{\mathbf{A}}^T \sqrt{\mathbf{A}} = \mathbf{A}$. Throughout this paper, “matrix square root” will refer to this definition (rather than $\sqrt{\mathbf{A}^2} = \mathbf{A}$), unless otherwise mentioned. Such a square root is not unique and is subject to the unitary freedom, i.e. if \mathbf{R} is a square root of \mathbf{A} , the set $\{\mathbf{Q}\mathbf{R} : \mathbf{Q}^T \mathbf{Q} = \mathbf{I}\}$ contains all square roots of \mathbf{A} . The choice of $\sqrt{\mathbf{A}}$ can be arbitrary but must be consistent. Moreover, we denote by $\mathbf{A}^{1/2}$ the unique SPD square root of \mathbf{A} .

For the optimization problem (1.1), we can always eliminate the inner product’s weights $\mathbf{\Gamma}_m$ and $\mathbf{\Gamma}_n$ by variable substitution. If we take $\mathbf{W} = \sqrt{\mathbf{\Gamma}_m} \mathbf{U} \sqrt{\mathbf{\Gamma}_n}^T$ and $\tilde{\mathbf{W}} = \sqrt{\mathbf{\Gamma}_m} \tilde{\mathbf{U}} \sqrt{\mathbf{\Gamma}_n}^T$, then (1.1) reduces to

$$\arg \min_{\tilde{\mathbf{W}} \in \mathcal{M}_{m \times n}: \tilde{\mathbf{W}}^T \tilde{\mathbf{W}} = \sqrt{\mathbf{\Gamma}_n} \tilde{\mathbf{P}} \sqrt{\mathbf{\Gamma}_n}^T} \|\tilde{\mathbf{W}} - \mathbf{W}\|_{\mathbf{F}}^2, \tag{1.3}$$

where $\|\mathbf{V}\|_{\mathbf{F}} = \text{tr}(\mathbf{V}^T \mathbf{V})$ is the unweighted Frobenius norm and $\sqrt{\mathbf{\Gamma}_n} \tilde{\mathbf{P}} \sqrt{\mathbf{\Gamma}_n}^T$ becomes the new target second-moment matrix. Therefore, without loss of generality, we will restrict ourselves to the simplified problem

$$\arg \min_{\tilde{\mathbf{U}} \in \mathcal{M}_{m \times n}: \tilde{\mathbf{U}}^T \tilde{\mathbf{U}} = \tilde{\mathbf{P}}} \|\tilde{\mathbf{U}} - \mathbf{U}\|_{\mathbf{F}}^2. \tag{1.4}$$

Note that we can also scale $\tilde{\mathbf{U}}$ to make $\tilde{\mathbf{P}} = \mathbf{I}$, but then $\mathbf{\Gamma}_n \neq \mathbf{I}$ in general. We choose to have $\mathbf{\Gamma}_n = \mathbf{I}$ and $\tilde{\mathbf{P}}$ arbitrary because it renders the problem’s structure symmetric and leads to richer and more insightful interpretations as will be seen in chapter 2. How the two ways of scaling are connected will be mentioned in section 2.3. Besides $\tilde{\mathbf{P}}$ being SPD, we further assume that $m \geq n$ and \mathbf{U} has full column rank, so $\mathbf{P} = \mathbf{U}^T \mathbf{U}$ is SPD. The rank-deficiency complications due to the violation of these conditions are addressed in section 2.6.

Minimum-correction second-moment matching has many applications. In the special case of $\tilde{\mathbf{P}} = \mathbf{I}$, the task reduces to minimum-correction orthonormalization. Many algorithms involve tracking a varying orthogonal matrix (Edelman et al., 1998; Rheinboldt, 1988), such as in optimization (Absil et al., 2009) or solving matrix differential equations with orthogonality constraints (Hairer et al.,

2006, sec. IV.9). There does exist orthogonality-preserving algorithms for some applications. For example, Wen and Yin (2013) studies an orthogonality-preserving curvilinear search algorithm for optimization on a Stiefel manifold, i.e. the manifold formed by orthogonal matrices of a particular size (Edelman et al., 1998). For time integration, Dieci et al. (1994) and Higham (1996) investigate orthogonality-preserving Gauss-Legendre Runge-Kutta schemes. However, such algorithms are typically constrained and the choices limited. Often, matrix update algorithms also prioritize other goals such as steepest descent optimization (Tagare, 2011) or accurate time integration (Koch and Lubich, 2007; Sapsis and Lermusiaux, 2009; Cheng et al., 2013; Ueckermann et al., 2013; Musharbash and Nobile, 2018). In such cases, after an orthogonal matrix is updated with a general-purpose algorithm, a deviation from orthogonality is usually incurred by the numerical discretization of the matrix’s continuous evolution on a Stiefel manifold. Therefore, we need to orthonormalize the matrix in such a way that this *entry-wise* continuous evolution is preserved (Feppon and Lermusiaux, 2018b; Govaerts and Werner, 1995). A natural idea is thus to find the closest point to the updated matrix on the Stiefel manifold (Dieci et al., 1994; Higham, 1997; Absil and Malick, 2012). In general, the procedure of pulling a point back to a manifold is referred to as a retraction (Absil et al., 2009). We will see a test case of this kind in chapter 5.

In the general case with $\tilde{\mathbf{P}} \neq \mathbf{I}$, one application is minimum-correction covariance matching. For example, in an ensemble Kalman filter (EnKF) (Evensen, 2003) for data assimilation, given the prior samples of a random vector in the rows of \mathbf{U} and some observation data, the goal is to obtain posterior samples $\tilde{\mathbf{U}}$ such that their empirical mean and covariance match those of the posterior distribution obtained by the Kalman update. A variant proposed by Ott et al. (2004) consists of updating the mean in the same way as EnKF but making the posterior sample variation $\tilde{\mathbf{U}}$ (with mean removed) as close as possible to the prior counterpart \mathbf{U} under the constraint that the empirical covariance $(1/m)\tilde{\mathbf{U}}^T\tilde{\mathbf{U}}$ must match the one obtained by the Gaussian Bayesian update. We will use such a test case in chapter 6.

Although the minimum-correction orthonormalization is relatively well understood with efficient algorithms developed, its general second-moment matching counterpart has not yet been studied and analyzed in a unified and systematic way. We aim to fill this gap in the present paper.

In chapter 2, we solve the optimization (1.4) analytically with a constructive proof (??), provide geometric intuition, solve rank deficient cases (2.6), and discuss the connections to other topics

including minimum congruence transform (section 2.2) and generalized polar decomposition (section 2.3). We also address the generalization to infinite dimensions in section 2.4 and connect to optimal transport in section 2.5. In chapter 3, we obtain numerical strategies for computing the optimal solution and review existing algorithms for polar decomposition and matrix square root that play an essential role in minimum-correction second-moment matching. We propose modification to stabilize two Newton iterations deemed unstable previously, which can now be used to efficiently compute both the orthogonal polar factor and the SPD square root. We compare algorithms in terms of accuracy, cost and robustness. Results are benchmarked using randomly generated matrices in chapter 4. In chapter 5, we show that these algorithms can be used in a subspace tracking method for stochastic ODEs and PDEs, the dynamically orthogonal equations (Sapsis and Lermusiaux, 2009; Feppon and Lermusiaux, 2018a), to maintain both the orthonormality and the time continuity of the base vectors. Furthermore, we show in chapter 6 how the algorithms for the general case can be applied to an ensemble-based filter for data assimilation, as introduced in Ott et al. (2004). Lastly, we conclude in chapter 7.

Chapter 2

Theory of Minimum-Correction Second-Moment Matching

2.1 Analytic Expression of the Global Optimizer

The constrained optimization (1.4) has a compact feasible region and a smooth objective function, so a global optimizer exists and is a critical point of the Lagrangian. Hence, we employ the approach of Lagrangian multipliers as in Ott et al. (2004). We complete a different proof, adopt more compact matrix notations, and provide new extensive discussions.

First, since the objective function of (1.4) can be rewritten as

$$\|\tilde{\mathbf{U}} - \mathbf{U}\|_{\text{F}}^2 = \text{tr}[(\tilde{\mathbf{U}} - \mathbf{U})^{\text{T}}(\tilde{\mathbf{U}} - \mathbf{U})] = \text{tr}\tilde{\mathbf{P}} + \text{tr}\mathbf{P} - 2\text{tr}(\mathbf{U}^{\text{T}}\tilde{\mathbf{U}}),$$

we have

$$\arg \min_{\tilde{\mathbf{U}} \in \mathcal{M}_{m \times n}: \tilde{\mathbf{U}}^{\text{T}}\tilde{\mathbf{U}} = \tilde{\mathbf{P}}} \|\tilde{\mathbf{U}} - \mathbf{U}\|_{\text{F}}^2 = \arg \min_{\tilde{\mathbf{U}} \in \mathcal{M}_{m \times n}: \tilde{\mathbf{U}}^{\text{T}}\tilde{\mathbf{U}} = \tilde{\mathbf{P}}} -2\text{tr}(\mathbf{U}^{\text{T}}\tilde{\mathbf{U}}). \quad (2.1)$$

The Lagrangian of this optimization is thus

$$L(\tilde{\mathbf{U}}, \mathbf{\Lambda}) = -2\text{tr}(\mathbf{U}^{\text{T}}\tilde{\mathbf{U}}) + \text{tr}(\mathbf{\Lambda}^{\text{T}}(\tilde{\mathbf{U}}^{\text{T}}\tilde{\mathbf{U}} - \tilde{\mathbf{P}})), \quad (2.2)$$

where $\Lambda \in \mathcal{M}_{n \times n}$ is the Lagrangian multiplier for the constraint $\tilde{U}^T \tilde{U} = \tilde{P}$. The global optimizer should be one of the critical points, so ¹

$$\nabla_{\tilde{U}} L|_{\tilde{U}_*, \Lambda_*} = -2\tilde{U} + 2\tilde{U}_* \Lambda_* = \mathbf{0}, \quad \nabla_{\Lambda} L|_{\tilde{U}_*, \Lambda_*} = \tilde{U}_*^T \tilde{U}_* - \tilde{P} = \mathbf{0}. \quad (2.3)$$

The first condition gives $\tilde{U} = \tilde{U}_* \Lambda_*$ and the second is simply the second-moment constraint on \tilde{U}_* . The symmetry of this constraint implies the symmetry of the Lagrangian multiplier Λ_* . Inserting these results into the objective function (2.1), we obtain

$$-2\text{tr}(\tilde{U}^T \tilde{U}_*) = -2\text{tr}(\Lambda \tilde{U}_*^T \tilde{U}_*) = -2\text{tr}(\Lambda_* \tilde{P}) = -2\text{tr}(\sqrt{\tilde{P}} \Lambda_* \sqrt{\tilde{P}}^T). \quad (2.4)$$

The reason for symmetrifying the matrix in the last step will soon become clear. Since

$$P = U^T U = \Lambda_* \tilde{U}_*^T \tilde{U}_* \Lambda_* = \Lambda_* \tilde{P} \Lambda_*,$$

the second-moment constraint reduces to $P = \Lambda_* \tilde{P} \Lambda_*$. To relate the objective function (2.4) to this constraint, we use a small trick:

$$\begin{aligned} (\sqrt{\tilde{P}} \Lambda_* \sqrt{\tilde{P}}^T)^2 &= (\sqrt{\tilde{P}} \Lambda_* \sqrt{\tilde{P}}^T)(\sqrt{\tilde{P}} \Lambda_* \sqrt{\tilde{P}}^T) \\ &= \sqrt{\tilde{P}} (\Lambda_* \tilde{P} \Lambda_*) \sqrt{\tilde{P}}^T \\ &= \sqrt{\tilde{P}} P \sqrt{\tilde{P}}^T. \end{aligned}$$

Therefore, $\sqrt{\tilde{P}} \Lambda_* \sqrt{\tilde{P}}^T$ must be a symmetric square root of $\sqrt{\tilde{P}} P \sqrt{\tilde{P}}^T$, which is not unique. This characterizes all the critical points of (1.4), among which is the desired global optimizer $\tilde{U}_* = U \Lambda_*^{-1}$.

Next, we identify the global optimizer by comparing the values of the objective function evaluated at these critical points. Since $\sqrt{\tilde{P}} \Lambda_* \sqrt{\tilde{P}}^T$ is symmetric, it has an eigendecomposition

$$\sqrt{\tilde{P}} \Lambda_* \sqrt{\tilde{P}}^T = V \text{diag}(\lambda_1, \dots, \lambda_n) V^T$$

¹See appendix A for how these matrix gradients are computed.

and its square $\sqrt{\tilde{\mathbf{P}}\mathbf{P}}\sqrt{\tilde{\mathbf{P}}}^T$ will have the corresponding eigendecomposition

$$\sqrt{\tilde{\mathbf{P}}\mathbf{P}}\sqrt{\tilde{\mathbf{P}}}^T = \mathbf{V} \text{diag}(\lambda_1^2, \dots, \lambda_n^2) \mathbf{V}^T. \quad (2.5)$$

If we denote the n positive eigenvalues of $\sqrt{\tilde{\mathbf{P}}\mathbf{P}}\sqrt{\tilde{\mathbf{P}}}^T$ by $\sigma_1, \dots, \sigma_n$, then without loss of generality, we must have $\sigma_i = \lambda_i^2$ and hence $\lambda_i = \pm\sqrt{\sigma_i}$. Therefore, the objective function (2.4) reduces to

$$-2\text{tr}(\mathbf{U}^T \tilde{\mathbf{U}}_*) = -2\text{tr}(\mathbf{V} \text{diag}(\lambda_1, \dots, \lambda_n) \mathbf{V}^T) = -2 \sum_{i=1}^n \pm\sqrt{\sigma_i} \geq -2 \sum_{i=1}^n \sqrt{\sigma_i}. \quad (2.6)$$

We can see that the lower bound is attained if and only if $\lambda_i = \sqrt{\sigma_i}$, $i = 1, \dots, n$, i.e. when $\sqrt{\tilde{\mathbf{P}}}\mathbf{\Lambda}_*\sqrt{\tilde{\mathbf{P}}}^T$ takes the unique SPD square root:

$$\sqrt{\tilde{\mathbf{P}}}\mathbf{\Lambda}_*\sqrt{\tilde{\mathbf{P}}}^T = (\sqrt{\tilde{\mathbf{P}}\mathbf{P}}\sqrt{\tilde{\mathbf{P}}}^T)^{1/2}. \quad (2.7)$$

Therefore, the global minimizer to (1.4) is $\tilde{\mathbf{U}}_* = \mathbf{U}\mathbf{\Lambda}_*^{-1}$ with

$$\mathbf{\Lambda}_*^{-1} = \sqrt{\tilde{\mathbf{P}}}^T (\sqrt{\tilde{\mathbf{P}}\mathbf{P}}\sqrt{\tilde{\mathbf{P}}}^T)^{-1/2} \sqrt{\tilde{\mathbf{P}}}. \quad (2.8)$$

Note that the particular choice of square root $\sqrt{\tilde{\mathbf{P}}}$ is irrelevant because both the spectrum of $\sqrt{\tilde{\mathbf{P}}\mathbf{P}}\sqrt{\tilde{\mathbf{P}}}^T$ and the global optimizer (2.8) are invariant under the unitary freedom $\sqrt{\tilde{\mathbf{P}}} \rightarrow \mathbf{Q}\sqrt{\tilde{\mathbf{P}}}$, due to the fact that $\mathbf{Q}\mathbf{A}^{1/2}\mathbf{Q}^T = (\mathbf{Q}\mathbf{A}\mathbf{Q}^T)^{1/2}$ for any SPD \mathbf{A} .

We remark that an alternative to using (2.4) is to use

$$-2\text{tr}(\mathbf{U}^T \tilde{\mathbf{U}}_*) = -2\text{tr}(\mathbf{U}^T \mathbf{U}\mathbf{\Lambda}_*^{-1}) = -2\text{tr}(\mathbf{P}\mathbf{\Lambda}_*^{-1}) = -2\text{tr}(\sqrt{\mathbf{P}}\mathbf{\Lambda}_*^{-1}\sqrt{\mathbf{P}}^T)$$

and correspondingly $\tilde{\mathbf{P}} = \mathbf{\Lambda}_*^{-1}\mathbf{P}\mathbf{\Lambda}_*^{-1}$. This simply switches the role of \mathbf{P} and $\tilde{\mathbf{P}}$ and replaces $\mathbf{\Lambda}_*$ with $\mathbf{\Lambda}_*^{-1}$. Therefore, this gives an alternative expression of $\mathbf{\Lambda}_*^{-1}$ obtained in (2.8), i.e.

$$\mathbf{\Lambda}_*^{-1} = \sqrt{\mathbf{P}}^{-1} (\sqrt{\mathbf{P}}\tilde{\mathbf{P}}\sqrt{\mathbf{P}}^T)^{1/2} \sqrt{\mathbf{P}}^{-T}. \quad (2.9)$$

We summarize all the results above in the following theorem.

Theorem 2.1 (Minimum-correction second-moment matching). *The global optimum of (1.4) is*

achieved by applying the unique $n \times n$ SPD linear transform to \mathbf{U} that matches the second-moment matrix of the result $\tilde{\mathbf{U}}$ to the target $\tilde{\mathbf{P}}$. More precisely,

$$\arg \min_{\tilde{\mathbf{U}} \in \mathcal{M}_{m \times n}: \tilde{\mathbf{U}}^T \tilde{\mathbf{U}} = \tilde{\mathbf{P}}} \|\tilde{\mathbf{U}} - \mathbf{U}\|_{\text{F}}^2 = \mathbf{U} \mathbf{A}_*, \quad (2.10)$$

$$\|\mathbf{U} \mathbf{A}_* - \mathbf{U}\|_{\text{F}}^2 = \text{tr}[\mathbf{P} + \tilde{\mathbf{P}} - 2(\sqrt{\tilde{\mathbf{P}}} \mathbf{P} \sqrt{\tilde{\mathbf{P}}}^T)^{1/2}] = \text{tr}[\mathbf{P} + \tilde{\mathbf{P}} - 2(\mathbf{P} \tilde{\mathbf{P}})^{1/2}] \quad (2.11)$$

² with the SPD linear transform \mathbf{A}_* being

$$\begin{aligned} \mathbf{A}_* &= \sqrt{\tilde{\mathbf{P}}}^T (\sqrt{\tilde{\mathbf{P}}} \mathbf{P} \sqrt{\tilde{\mathbf{P}}}^T)^{-1/2} \sqrt{\tilde{\mathbf{P}}} \\ &= \sqrt{\mathbf{P}}^{-1} (\sqrt{\mathbf{P}} \tilde{\mathbf{P}} \sqrt{\mathbf{P}}^T)^{1/2} \sqrt{\mathbf{P}}^{-T}. \end{aligned} \quad (2.12)$$

Note that when $\tilde{\mathbf{P}} = \mathbf{P}$, we always have $\mathbf{A}_* = \mathbf{I}$, which correctly gives the null correction as the optimizer in this trivial case. Moreover, $\tilde{\mathbf{U}}_* = \mathbf{U} \mathbf{A}_*$ implies $\text{Col}(\tilde{\mathbf{U}}_*) = \text{Col}(\mathbf{U})$, which is particularly important for covariance matching because a zero-mean \mathbf{U} (i.e. $\mathbf{e}^T \mathbf{U} = \mathbf{0}$, where $\mathbf{e}^T = [1, \dots, 1]$) will yield a zero-mean $\tilde{\mathbf{U}}_*$.

Since $\mathcal{M}_{m \times n}$ can be quipped with the inner product $\langle \mathbf{U}, \mathbf{V} \rangle_{\text{F}} \triangleq \text{tr}(\mathbf{U}^T \mathbf{V})$ which induces the Frobenius norm $\|\cdot\|_{\text{F}}$, theorem 2.1 indeed characterizes the orthogonal projection $\mathcal{P}_{\tilde{\mathbf{P}}}$ of any $\mathbf{U} \in \mathcal{M}_{m \times n}$ onto the $(mn - \frac{1}{2}n(n+1))$ dimensional sub-manifold $\mathcal{S}_{\tilde{\mathbf{P}}} = \{\tilde{\mathbf{U}} \in \mathcal{M}_{m \times n} : \tilde{\mathbf{U}}^T \tilde{\mathbf{U}} = \tilde{\mathbf{P}}\}$ as the unique $n \times n$ SPD linear transform that maps \mathbf{U} into the sub-manifold.

Moreover, if we invert the first expression in (2.12), we obtain

$$\mathbf{A}_*^{-1} = \sqrt{\tilde{\mathbf{P}}}^{-1} (\sqrt{\tilde{\mathbf{P}}} \mathbf{P} \sqrt{\tilde{\mathbf{P}}}^T)^{1/2} \sqrt{\tilde{\mathbf{P}}}^{-T},$$

which is nothing but the second expression in (2.12) with $\tilde{\mathbf{P}}$ and \mathbf{P} switched. This implies that

² For square \mathbf{A} diagonalizable with nonnegative eigenvalues and EVD $\mathbf{A} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^{-1}$, we define $\mathbf{A}^{1/2} \triangleq \mathbf{V} \mathbf{\Lambda}^{1/2} \mathbf{V}^{-1}$. Since $\mathbf{P} \tilde{\mathbf{P}} = \sqrt{\tilde{\mathbf{P}}}^{-1} (\sqrt{\tilde{\mathbf{P}}} \mathbf{P} \sqrt{\tilde{\mathbf{P}}}^T) \sqrt{\tilde{\mathbf{P}}}$ is similar to $\sqrt{\tilde{\mathbf{P}}} \mathbf{P} \sqrt{\tilde{\mathbf{P}}}^T$, $\mathbf{P} \tilde{\mathbf{P}}$, though not symmetric in general, is diagonalizable with positive eigenvalues and EVD $\mathbf{P} \tilde{\mathbf{P}} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^{-1}$. Therefore, $\text{tr}((\mathbf{P} \tilde{\mathbf{P}})^{1/2}) = \sum_{i=1}^n \sqrt{\lambda_i} = \text{tr}((\sqrt{\tilde{\mathbf{P}}} \mathbf{P} \sqrt{\tilde{\mathbf{P}}}^T)^{1/2})$.

$\tilde{U}A_*^{-1} = U$ projects \tilde{U} orthogonally onto \mathcal{S}_P , i.e.

$$\arg \min_{U \in \mathcal{M}_{m \times n}: U^T U = P} \|U - \tilde{U}\|_F^2 = \tilde{U}A_*^{-1} \quad (2.13)$$

with A_* given by (2.12). Therefore, the orthogonal projection \mathcal{P}_P and $\mathcal{P}_{\tilde{P}}$ specify a bijection between \mathcal{S}_P and $\mathcal{S}_{\tilde{P}}$ such that

$$\mathcal{P}_P \circ \mathcal{P}_{\tilde{P}}|_{\mathcal{S}_P} = \text{Id}|_{\mathcal{S}_P}, \quad \mathcal{P}_{\tilde{P}} \circ \mathcal{P}_P|_{\mathcal{S}_{\tilde{P}}} = \text{Id}|_{\mathcal{S}_{\tilde{P}}}. \quad (2.14)$$

In addition, since A_* depends on P and \tilde{P} only, both $\mathcal{P}_{\tilde{P}}|_{\mathcal{S}_P}$ and $\mathcal{P}_P|_{\mathcal{S}_{\tilde{P}}}$ are characterized by a constant SPD linear transform on \mathbb{R}^n .

Furthermore, the simple case of $n = 1$ helps us build a geometric intuition of the above results, as illustrated in fig. 2-1. When $n = 1$, P and \tilde{P} reduce to a positive real number, while \mathcal{S}_P and $\mathcal{S}_{\tilde{P}}$ reduce to two concentric $(m - 1)$ -spheres (or two $(m - 1)$ -ellipsoids with the same eccentricity if $\Gamma_m \neq \lambda I$). Hence, it is straightforward to see that the rays from the origin provide a bijection between the two spheres. This mapping is equivalent to scaling one sphere by a constant ratio. As fig. 2-1 shows, the points on the two manifolds can be paired up such that in each point pair, the blue point is the closet point on the blue manifold to the red one and vice versa.

However, the case of $n = 1$ also over-simplifies some aspects of the problem and could be misleading. For example, fig. 2-1 could create the illusion that the projection operator is transitive, i.e. projecting from \mathcal{S}_{P_1} to \mathcal{S}_{P_3} is equivalent to first projecting from \mathcal{S}_{P_1} to some intermediate \mathcal{S}_{P_2} and then from \mathcal{S}_{P_2} to \mathcal{S}_{P_3} . We thus emphasize that in general,

$$\mathcal{P}_{P_3} \circ \mathcal{P}_{P_2}|_{\mathcal{S}_{P_1}} \neq \mathcal{P}_{P_3}|_{\mathcal{S}_{P_1}}, \quad A_*^{1 \rightarrow 2} A_*^{2 \rightarrow 3} \neq A_*^{1 \rightarrow 3}, \quad (2.15)$$

since $A_*^{1 \rightarrow 3}$ is symmetric but $A_*^{1 \rightarrow 2} A_*^{2 \rightarrow 3}$ is not in general. This is especially clear when $P_2 = I$, which implies $A_*^{1 \rightarrow 2} = P_1^{-1/2}$ and $A_*^{2 \rightarrow 3} = P_3^{1/2}$ according to (2.12). Hence, whenever P_1 and P_3 are not commutative, $A_*^{1 \rightarrow 2} A_*^{2 \rightarrow 3}$ is not symmetric.

Finally, with (1.3), we write down the global optimizer for the general setting (1.1).

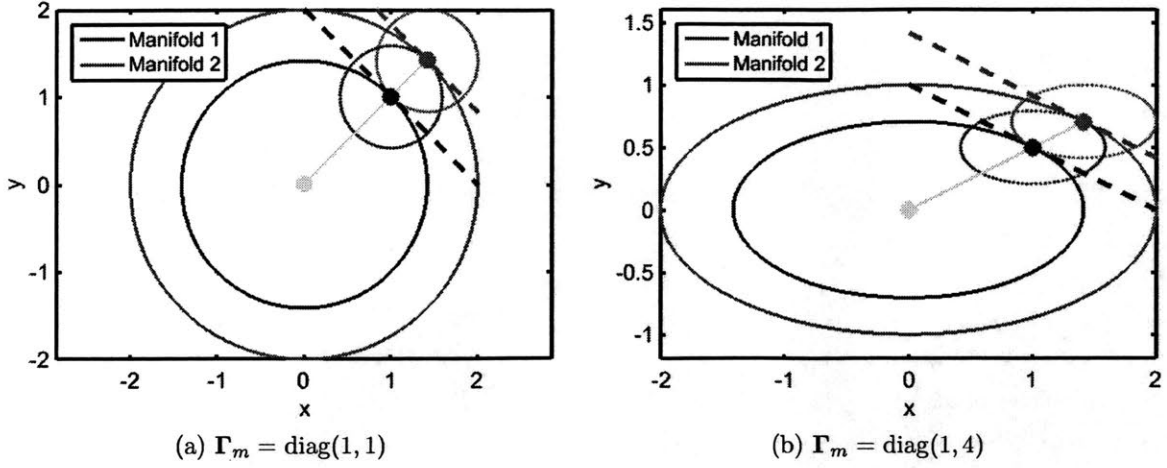


Figure 2-1: A geometric picture of theorem 2.1 for $m = 2$ and $n = 1$.

Corollary 2.2 (General setting with inner product weighted by Γ_m and Γ_n).

$$\arg \min_{\tilde{U} \in \mathcal{M}_{m \times n}: \tilde{U}^T \Gamma_m \tilde{U} = \tilde{P}} \|\tilde{U} - U\|_{F, \Gamma_m, \Gamma_n}^2 = U(\sqrt{\Gamma_n}^T A_* \sqrt{\Gamma_n}^{-T}) \quad (2.16)$$

$$\|\tilde{U}_* - U\|_F^2 = \text{tr}[P_\Gamma + \tilde{P}_\Gamma - 2(P_\Gamma \tilde{P}_\Gamma)^{1/2}] \quad (2.17)$$

$$\begin{aligned} A_* &= \sqrt{\tilde{P}_\Gamma}^T (\sqrt{\tilde{P}_\Gamma} P_\Gamma \sqrt{\tilde{P}_\Gamma}^T)^{-1/2} \sqrt{\tilde{P}_\Gamma} \\ &= \sqrt{P_\Gamma}^{-1} (\sqrt{P_\Gamma} \tilde{P}_\Gamma \sqrt{P_\Gamma}^T)^{1/2} \sqrt{P_\Gamma}^{-T}. \end{aligned} \quad (2.18)$$

$$\tilde{P}_\Gamma = \sqrt{\Gamma_n} \tilde{P} \sqrt{\Gamma_n}^T, \quad P_\Gamma = \sqrt{\Gamma_n} P \sqrt{\Gamma_n}^T. \quad (2.19)$$

We have shown previously that the choice of square root $\sqrt{\tilde{P}_\Gamma}$ and $\sqrt{P_\Gamma}$ in (2.18) is irrelevant. Here we can show that the choice of square root $\sqrt{\Gamma_n}$ is also irrelevant. Under the unitary freedom $\sqrt{\Gamma_n} \rightarrow Q\sqrt{\Gamma_n}$ and thus $P_\Gamma \rightarrow QP_\Gamma Q^T$, we can choose $\sqrt{QP_\Gamma Q^T}$ to be $Q\sqrt{P_\Gamma}Q^T$. Therefore, the Q 's will pair up with the Q^T 's and cancel each other, which renders the linear transform $\sqrt{\Gamma_n}^T A_* \sqrt{\Gamma_n}^{-T}$ invariant. However, the choice of the inner product weights Γ_m and Γ_n themselves does affect the optimum, although they play quite different roles.

Γ_m defines the inner product on \mathbb{R}^m for computing the second-moment matrices. This inner product

usually has physical meaning and is commonly determined by information from the underlying problem. For example, if each column of \mathbf{U} is a discrete representation of a function $f(x)$ over some domain Ω , then $\mathbf{\Gamma}_m$ could be chosen to make $\mathbf{u}^T \mathbf{\Gamma}_m \mathbf{v}$ an approximation to inner product $\int_{\Omega} u(x)v(x)dx$ and the second-moment matrix becomes a Gram matrix. If each row of \mathbf{U} is a sample of a zero-mean random vector, then $\mathbf{\Gamma}_m = \frac{1}{m} \text{diag}(1, \dots, 1)$ will make $\mathbf{u}^T \mathbf{\Gamma}_m \mathbf{v}$ the empirical covariance matrix. Note that \mathbf{A}_* does not depend on $\mathbf{\Gamma}_m$ explicitly, but only through how $\tilde{\mathbf{P}}$ and \mathbf{P} are computed based on $\mathbf{\Gamma}_m$.

In contrast, $\mathbf{\Gamma}_n$ is more of a tunable numerical parameter. It controls how we want to weight the columns of \mathbf{U} differently so that when we evaluate the closeness between $\tilde{\mathbf{U}}$ and \mathbf{U} , we give priority to some columns over others. Such differentiation could be advantageous if, for example, some of the columns of \mathbf{U} contain more reliable information (e.g. due to smaller numerical errors) than others, or the columns are different by several orders of magnitude and the small columns barely contribute to the closeness metric. If the magnitude is of concern, a natural choice of a scaling matrix is either $\mathbf{\Gamma}_n = \mathbf{P}^{-1}$ or $\mathbf{\Gamma}_n = \tilde{\mathbf{P}}^{-1}$:

$$\begin{aligned}\mathbf{\Gamma}_n = \mathbf{P}^{-1} &\implies \sqrt{\mathbf{\Gamma}_n}^T \mathbf{A}_* \sqrt{\mathbf{\Gamma}_n}^{-T} = \sqrt{\mathbf{P}}^{-1} (\sqrt{\mathbf{P}}^{-T} \tilde{\mathbf{P}} \sqrt{\mathbf{P}}^{-1})^{1/2} \sqrt{\mathbf{P}} \\ \mathbf{\Gamma}_n = \tilde{\mathbf{P}}^{-1} &\implies \sqrt{\mathbf{\Gamma}_n}^T \mathbf{A}_* \sqrt{\mathbf{\Gamma}_n}^{-T} = \sqrt{\tilde{\mathbf{P}}}^{-1} (\sqrt{\tilde{\mathbf{P}}}^{-T} \mathbf{P} \sqrt{\tilde{\mathbf{P}}}^{-1})^{-1/2} \sqrt{\tilde{\mathbf{P}}}.\end{aligned}$$

Surprisingly, as Ott et al. (2004) points out, these two choices of $\mathbf{\Gamma}_n$ lead to the same optimal mapping. By the properties (2.13) and (2.14) of the projection operators, we can quickly show the equivalence between these two expressions by noticing that either one can be obtained by inverting the other and switching \mathbf{P} with $\tilde{\mathbf{P}}$. Last but not least, note that the optimal mapping $\sqrt{\mathbf{\Gamma}_n}^T \mathbf{A}_* \sqrt{\mathbf{\Gamma}_n}^{-T}$ explicitly depends on $\mathbf{\Gamma}_n$ and when $\mathbf{\Gamma}_n$ is not a scalar matrix, the optimal mapping is in general not symmetric.

2.2 Minimum Linear Operation for Second-Moment Matching and Minimum Congruence Transform

In the optimization problem (1.4), the candidate $\tilde{\mathbf{U}}$ can come from all of $\mathcal{M}_{m \times n}$, but the global optimizer $\mathbf{U} \mathbf{A}_*$ turns out to have the property of sharing the column space of \mathbf{U} . This implies that

if we restrict the candidate set to the n^2 -dimensional subspace $\{\tilde{U} \in \mathcal{M}_{m \times n} : \text{Im}(\tilde{U}) \subset \text{Im}(U)\} = \{\tilde{U} = UA : A \in \mathcal{M}_{n \times n}\}$ of the mn -dimensional $\mathcal{M}_{m \times n}$, the global optimizer is still given by (2.12). This is equivalent to requiring the second-moment matching to be achieved by an n -by- n linear transform. With this restriction, the objective function in (1.4) becomes

$$\begin{aligned} \|\tilde{U} - U\|_{\text{F}}^2 &= \|UA - U\|_{\text{F}}^2 \\ &= \text{tr}[(A - I)^T U^T U (A - I)] \\ &= \text{tr}((A - I)^T P (A - I)) \\ &= \|A - I\|_{\text{F}, P}^2, \end{aligned} \tag{2.20}$$

where the minimum correction of U translates to the minimum operation of A , quantified by being the closest linear transform to the identity mapping. Besides, the second-moment matching constraint is now on A

$$\tilde{P} = \tilde{U}^T \tilde{U} = A^T U^T U A = A^T P A.$$

Therefore, theorem 2.1 implies the following corollary.

Corollary 2.3 (Minimum linear operation for second-moment matching and minimum congruence transform).

$$\arg \min_{A \in \mathcal{M}_{n \times n} : A^T P A = \tilde{P}} \|A - I\|_{\text{F}, P}^2 = A_* \tag{2.21}$$

with A_* given by (2.12) for any SPD \tilde{P} and P in $\mathcal{M}_{n \times n}$.

Note that this optimization problem is related to U only through its second-moment matrix P . We can indeed interpret it in a way that does not involve U . Since $A^T P A = \tilde{P}$ indicates that \tilde{P} is congruent with P through the congruence transform A , (2.21) can be viewed as finding the minimum congruence transform between two SPD matrices. Here “minimum” again refers to the minimum action of an operator. The symmetry result (2.13) implies that if P and \tilde{P} are switched in (2.21), the global optimizer simply becomes A_*^{-1} .

The connection of (2.21) to (1.4) through (2.20) reveals a natural and convenient choice of norm for measuring the distance between A and I . The key here is that the Frobenius norm weighted by

P renders the objective function linear:

$$\begin{aligned}\|A - I\|_{F,P}^2 &= \text{tr}(A^T P A) - \text{tr}(P A) - \text{tr}(A^T P) + \text{tr}(P) \\ &= \text{tr}(\tilde{P}) - 2\text{tr}(P A) + \text{tr}(P),\end{aligned}$$

which enables solving the problem analytically. If we use $\|\cdot\|_F$ or even $\|\cdot\|_{F,\tilde{P}}$ instead, the quadratic term above will not become a constant and will have to remain in the objective function. Note that although the unweighted Frobenius norm has the property $\|A\|_F = \|A^T\|_F$, a weighted one does not in general. First, only for square A are both $\|A\|_{F,P}$ and $\|A^T\|_{F,P}$ defined. Second, unless A is normal or A commutes with P , typically $\|A\|_{F,P} \neq \|A^T\|_{F,P}$.

2.3 Minimum-Correction Orthonormalization and Polar Decomposition with Generalization

An important and familiar special case of the problem (1.4) is when $\tilde{P} = I$ and the task reduces to minimum correction orthonormalization, i.e. finding the closest orthogonal matrix \tilde{U} to an arbitrary $U \in \mathcal{M}_{m \times n}$ with full column rank. In this case, the candidate solutions form the set of all m -by- n orthogonal matrices $\mathcal{S}_I = \{V \in \mathcal{M}_{m \times n} : V^T V = I\}$, which is known as the Stiefel manifold. Therefore, theorem 2.1 reduces to the well-known fact that a matrix can be orthogonally projected onto the Stiefel manifold through its polar decomposition (Absil and Malick, 2012), summarized in the following corollary.

Corollary 2.4 (Orthogonal projection onto the Stiefel manifold by polar decomposition). *Given $U \in \mathcal{M}_{m \times n}$ with full column rank and $P = U^T U$, we have*

$$\tilde{U}_* = \arg \min_{\tilde{U} \in \mathcal{S}_I} \|\tilde{U} - U\|_F^2 = U P^{-1/2}, \quad (2.22)$$

where $U = \tilde{U}_* P^{1/2}$ gives the unique polar decomposition of U .

Note that corollary 2.2 implies that the above result readily generalizes to the case with a Frobenius

norm weighted by $\mathbf{\Gamma}_m$ and $\mathbf{\Gamma}_n$; the polar decomposition then becomes the weighted polar decomposition investigated in Yang and Li (2008). However, our results will likely not generalize to the generalized polar decomposition studied in Higham et al. (2010), where the inner products on \mathbb{R}^m and \mathbb{R}^n are relaxed to an ortho-symmetric pair of scalar products, where $\mathbf{\Gamma}_m$ and $\mathbf{\Gamma}_n$ are allowed to be indefinite or even non-symmetric. An indefinite $\mathbf{\Gamma}$ renders the quadratic form $\mathbf{x}^T \mathbf{\Gamma} \mathbf{x}$ non-convex and makes $\mathbf{x} = 0$ a saddle point rather than a global minimizer, which likely makes the generalized polar decomposition defy interpretation from a variational perspective.

On the other hand, our results with $\tilde{\mathbf{P}} \neq \mathbf{I}$ in theorem 2.1 do provide another way to generalize the polar decomposition.

Corollary 2.5. *Given $\mathbf{U} \in \mathcal{M}_{m \times n}$ with full column rank and $\tilde{\mathbf{P}} \in \mathcal{M}_{n \times n}$ that is SPD, we have a unique polar-like decomposition*

$$\mathbf{U} = \tilde{\mathbf{U}} \mathbf{\Lambda} \tag{2.23}$$

such that $\tilde{\mathbf{U}} \in \mathcal{M}_{m \times n}$, $\tilde{\mathbf{U}}^T \tilde{\mathbf{U}} = \tilde{\mathbf{P}}$ and $\mathbf{\Lambda} \in \mathcal{M}_{n \times n}$ is SPD. Moreover, this decomposition orthogonally projects \mathbf{U} onto the sub-manifold $\mathcal{S}_{\tilde{\mathbf{P}}}$, and $\tilde{\mathbf{U}}$ is the closest point to \mathbf{U} on $\mathcal{S}_{\tilde{\mathbf{P}}}$.

This corollary enriches the polar decomposition with the symmetry between \mathbf{U} and $\tilde{\mathbf{U}}$. Note that this symmetry is preserved by our choice of scaling that simplifies (1.1) to (1.4). Should we scale $\tilde{\mathbf{P}}$ rather than $\mathbf{\Gamma}_n$ to \mathbf{I} , we would instead obtain the weighted polar decomposition in Yang and Li (2008).

2.4 Generalization from \mathbb{R}^m to a Hilbert Space

In theorem 2.1, each column of \mathbf{U} or $\tilde{\mathbf{U}}$ is an element in \mathbb{R}^m and the second-moment matrices are computed based on an inner product on \mathbb{R}^m . A natural question is whether theorem 2.1 still holds when \mathbb{R}^m is replaced by an infinite-dimensional Hilbert space \mathcal{H} equipped with an inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$.

More precisely, assume $u_1, \dots, u_n \in \mathcal{H}$ are linearly independent and form an n -vector of \mathcal{H} denoted by $\mathbf{u} = [u_1, \dots, u_n]^T \in \mathcal{H}^n$. Then we can compute the second-moment matrix (a.k.a. the Gram

matrix) of \mathbf{u} as $\mathbf{P} = \text{Gr}(\mathbf{u}, \mathbf{u})$, where

$$\text{Gr}(\mathbf{u}, \mathbf{v}) \triangleq \langle \mathbf{u}, \mathbf{v}^T \rangle_{\mathcal{H}} \in \mathcal{M}_{n \times n} \quad (2.24)$$

i.e. with the (i, j) -th element of $\text{Gr}(\mathbf{u}, \mathbf{v})$ being $\langle u_i, v_j \rangle_{\mathcal{H}}$. If \mathcal{H}^n is equipped with an inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}^n}$ defined as

$$\langle \mathbf{u}, \mathbf{v} \rangle_{\mathcal{H}^n} \triangleq \text{tr} \left(\sqrt{\mathbf{\Gamma}_n} \text{Gr}(\mathbf{u}, \mathbf{v}) \sqrt{\mathbf{\Gamma}_n}^T \right) \quad (2.25)$$

with again an SPD weight matrix $\mathbf{\Gamma}_n$, we want to solve the following minimum-correction second-moment matching problem

$$\arg \min_{\tilde{\mathbf{u}} \in \mathcal{H}^n : \text{Gr}(\tilde{\mathbf{u}}, \tilde{\mathbf{u}}) = \tilde{\mathbf{P}}} \|\tilde{\mathbf{u}} - \mathbf{u}\|_{\mathcal{H}^n}^2 \quad (2.26)$$

for a given SPD $\tilde{\mathbf{P}} \in \mathcal{M}_{n \times n}$. Here the norm is induced by the inner product (2.25). As a common scenario, if $\mathcal{H} = L_2(\Omega)$, which is the space of square-integrable real functions on some measure space (Ω, \mathcal{A}, m) equipped with inner product $\langle u(\cdot), v(\cdot) \rangle_{\mathcal{H}} = \int_{\Omega} u(\omega)v(\omega)dm(\omega)$, then the inner product (2.25) on \mathcal{H}^n can be rewritten as

$$\langle \mathbf{u}, \mathbf{v} \rangle_{\mathcal{H}^n} = \int_{\Omega} \mathbf{u}(\omega)^T \mathbf{\Gamma}_n \mathbf{v}(\omega) dm(\omega). \quad (2.27)$$

The definitions (2.25) and (2.27) are simply the Hilbert space counterparts of the two expressions in (1.2). Since we can get rid of $\mathbf{\Gamma}_n$ by scaling the elements in \mathcal{H}^n by $\sqrt{\mathbf{\Gamma}_n}$ as before, we will assume $\mathbf{\Gamma}_n = \mathbf{I}$ hereafter.

theorem 2.1 does not apply to (2.26), but as in section 2.2, if we restrict the candidate set to $\{\tilde{\mathbf{u}} = \mathbf{A}\mathbf{u} : \mathbf{A} \in \mathcal{M}_{n \times n}\}$, the Hilbert space complication will be confined to computing the Gram matrices only. Since

$$\text{Gr}(\mathbf{A}\mathbf{u}, \mathbf{B}\mathbf{v}) = \langle \mathbf{A}\mathbf{u}, (\mathbf{B}\mathbf{v})^T \rangle_{\mathcal{H}} = \mathbf{A} \langle \mathbf{u}, \mathbf{v}^T \rangle_{\mathcal{H}} \mathbf{B}^T = \mathbf{A} \text{Gr}(\mathbf{u}, \mathbf{v}) \mathbf{B}^T$$

due to the bilinearity of an inner product, we have

$$\|\tilde{\mathbf{u}} - \mathbf{u}\|_{\mathcal{H}^n}^2 = \|(\mathbf{A} - \mathbf{I})\mathbf{u}\|_{\mathcal{H}^n}^2 = \text{tr} \left((\mathbf{A} - \mathbf{I}) \text{Gr}(\mathbf{u}, \mathbf{u}) (\mathbf{A} - \mathbf{I})^T \right) = \|\mathbf{A} - \mathbf{I}\|_{\mathbf{F}, \mathbf{P}}^2.$$

Hence (2.26) reduces to (2.21) and corollary 2.3 applies.

To establish the analog of theorem 2.1 for (2.26), we need some techniques from the calculus of variation to generalize optimization in \mathbb{R}^m to that in a Hilbert space.

Theorem 2.6 (Minimum-correction 2nd-moment matching on a Hilbert space). *Given $\mathbf{u} \in \mathcal{H}^n$ whose entries are linearly independent and $\tilde{\mathbf{P}}$ which is SPD, we have*

$$\arg \min_{\tilde{\mathbf{u}} \in \mathcal{H}^n: \text{Gr}(\tilde{\mathbf{u}}, \tilde{\mathbf{u}}) = \tilde{\mathbf{P}}} \|\tilde{\mathbf{u}} - \mathbf{u}\|_{\mathcal{H}^n}^2 = \mathbf{A}_* \mathbf{u} \quad (2.28)$$

with \mathbf{A}_* given by (2.12).

Proof. Compared to the case in \mathbb{R}^m , although $\mathcal{S}_{\tilde{\mathbf{P}}} = \{\tilde{\mathbf{u}} \in \mathcal{H}^n : \text{Gr}(\tilde{\mathbf{u}}, \tilde{\mathbf{u}}) = \tilde{\mathbf{P}}\}$ is still closed and bounded ($\|\tilde{\mathbf{u}}\|_{\mathcal{H}^n}^2 \equiv \text{tr}(\tilde{\mathbf{P}})$ on $\mathcal{S}_{\tilde{\mathbf{P}}}$), $\mathcal{S}_{\tilde{\mathbf{P}}}$ is no longer compact due to the infinite dimensionality. Therefore, although the objective function is continuous and bounded from below, its infimum may not be attainable. In the following, we first assume that a global minimizer $\tilde{\mathbf{u}}_*$ exists and derive an analytic expression for it. After that, we will show that this expression, which is well-defined whether or not a global minimum exists, is indeed the unique global minimizer to (2.26).

Here the objective function is

$$F(\tilde{\mathbf{u}}) = \|\tilde{\mathbf{u}} - \mathbf{u}\|_{\mathcal{H}^n}^2 = \text{tr}(\text{Gr}(\tilde{\mathbf{u}} - \mathbf{u}, \tilde{\mathbf{u}} - \mathbf{u})) = \text{tr}(\tilde{\mathbf{P}}) + \text{tr}(\mathbf{P}) - 2\text{tr}(\text{Gr}(\tilde{\mathbf{u}}, \mathbf{u})),$$

so its first-order variation is

$$\delta F(\tilde{\mathbf{u}}, \delta \tilde{\mathbf{u}}) = -2\text{tr}(\text{Gr}(\delta \tilde{\mathbf{u}}, \mathbf{u})).$$

The constraint on $\tilde{\mathbf{u}}$ is $\text{Gr}(\tilde{\mathbf{u}}, \tilde{\mathbf{u}}) = \tilde{\mathbf{P}}$, so the constraint on $\delta \tilde{\mathbf{u}}$ is

$$\text{Gr}(\delta \tilde{\mathbf{u}}, \tilde{\mathbf{u}}) + \text{Gr}(\tilde{\mathbf{u}}, \delta \tilde{\mathbf{u}}) = 0, \quad (2.29)$$

which is simply saying that $\text{Gr}(\delta \tilde{\mathbf{u}}, \tilde{\mathbf{u}})$ is anti-symmetric. If $\tilde{\mathbf{u}}$ is a local minimizer to (2.28), we must have $\delta F(\tilde{\mathbf{u}}, \delta \tilde{\mathbf{u}}) \equiv 0$ for any $\delta \tilde{\mathbf{u}} \in \mathcal{H}^n$ that satisfies (2.29).

Now we prove by contradiction that each entry of \mathbf{u} must be a linear combination of the entries of a local minimizer $\tilde{\mathbf{u}}$, which means $\mathbf{u} = \mathbf{A}\tilde{\mathbf{u}}$ for some $\mathbf{A} \in \mathcal{M}_{n \times n}$. If this was not the case, without loss of generality, we can assume $u_1 \notin \text{span}\{\tilde{u}_1, \dots, \tilde{u}_n\} \subset \mathcal{H}$. Therefore, the projection of u_1 onto

$\text{span}\{\tilde{u}_1, \dots, \tilde{u}_n\}^\perp$ must be nonzero and we denote it by v . Now we construct $\delta\tilde{\mathbf{u}} = [v, 0, \dots, 0]^\text{T}$. We can see that $\text{Gr}(\delta\tilde{\mathbf{u}}, \tilde{\mathbf{u}}) = 0$ because $\langle v, \tilde{u}_i \rangle_{\mathcal{H}} = 0$, $i = 1, \dots, n$. Hence, (2.29) is satisfied. On the other hand, since $\langle v, u_1 \rangle_{\mathcal{H}} = \|v\|_{\mathcal{H}}^2 > 0$, we have

$$\delta F(\tilde{\mathbf{u}}, \delta\tilde{\mathbf{u}}) = -2\text{tr}(\text{Gr}(\delta\tilde{\mathbf{u}}, \mathbf{u})) = -2 \sum_{i=1}^n \langle \delta\tilde{u}_i, u_i \rangle_{\mathcal{H}} = -2 \langle v, u_1 \rangle_{\mathcal{H}} < 0,$$

which contradicts the assumption that $\tilde{\mathbf{u}}$ is a local minimizer. Therefore, any local minimizer $\tilde{\mathbf{u}}$ must satisfy $\mathbf{u} = \mathbf{A}\tilde{\mathbf{u}}$ for some $\mathbf{A} \in \mathcal{M}_{n \times n}$. Since the entries of \mathbf{u} are linearly independent, \mathbf{A} must be invertible and thus $\tilde{\mathbf{u}} = \mathbf{A}^{-1}\mathbf{u}$.

Since we have already shown that with the candidate set restricted to $\{\tilde{\mathbf{u}} = \mathbf{A}\mathbf{u}\}$, corollary 2.3 applies. This completes the proof of (2.28) when a global minimum exists.

Note that no matter whether a global minimum exists or not, we can always construct $\tilde{\mathbf{u}}_* = \mathbf{A}_*\mathbf{u} \in \mathcal{S}_{\tilde{\mathbf{P}}}$ with \mathbf{A}_* given by (2.12). Therefore, if we can show $\|\tilde{\mathbf{u}}_* - \mathbf{u}\|_{\mathcal{H}^n}$ is indeed a lower bound of $\|\tilde{\mathbf{u}} - \mathbf{u}\|_{\mathcal{H}^n}$ on $\mathcal{S}_{\tilde{\mathbf{P}}}$, we can prove the existence and uniqueness of a global minimizer.

Since $F(\tilde{\mathbf{u}}) - \text{tr}(\tilde{\mathbf{P}}) - \text{tr}(\mathbf{P}) = -2\text{tr}(\text{Gr}(\tilde{\mathbf{u}}, \mathbf{u})) = -2 \langle \tilde{\mathbf{u}}, \mathbf{u} \rangle_{\mathcal{H}^n}$, we want to show $\langle \tilde{\mathbf{u}}, \mathbf{u} \rangle_{\mathcal{H}^n} \leq \langle \tilde{\mathbf{u}}_*, \mathbf{u} \rangle_{\mathcal{H}^n}$ for any $\tilde{\mathbf{u}} \in \mathcal{S}_{\tilde{\mathbf{P}}}$ by the Cauchy-Schwarz inequality. Since $\tilde{\mathbf{u}}_*$ is not parallel to \mathbf{u} in general, we need two linear transforms $\tilde{\mathbf{u}} = \mathbf{A}_*^{1/2}\tilde{\mathbf{v}}$ and $\mathbf{u} = \mathbf{A}_*^{-1/2}\mathbf{v}$ to make sure that the Cauchy-Schwarz inequality attains its equal sign when $\tilde{\mathbf{u}} = \mathbf{A}_*\mathbf{u}$. Therefore, we have

$$\langle \tilde{\mathbf{u}}, \mathbf{u} \rangle_{\mathcal{H}^n} = \text{tr} \left(\mathbf{A}_*^{1/2} \langle \tilde{\mathbf{v}}, \mathbf{v}^\text{T} \rangle_{\mathcal{H}} \mathbf{A}_*^{-1/2} \right) = \text{tr} \left(\langle \tilde{\mathbf{v}}, \mathbf{v}^\text{T} \rangle_{\mathcal{H}} \right) = \langle \tilde{\mathbf{v}}, \mathbf{v} \rangle_{\mathcal{H}^n} \leq \|\tilde{\mathbf{v}}\|_{\mathcal{H}^n} \|\mathbf{v}\|_{\mathcal{H}^n},$$

where the last step is the Cauchy-Schwarz inequality. On the other hand, we have

$$\begin{aligned} \|\mathbf{v}\|_{\mathcal{H}^n}^2 &= \left\langle \mathbf{A}_*^{1/2}\mathbf{u}, \mathbf{A}_*^{1/2}\mathbf{u} \right\rangle_{\mathcal{H}^n} = \text{tr} \left(\mathbf{A}_*^{1/2} \langle \tilde{\mathbf{u}}, \mathbf{u}^\text{T} \rangle_{\mathcal{H}} \mathbf{A}_*^{1/2} \right) = \text{tr}(\mathbf{A}_*\mathbf{P}), \\ \|\tilde{\mathbf{v}}\|_{\mathcal{H}^n}^2 &= \left\langle \mathbf{A}_*^{-1/2}\tilde{\mathbf{u}}, \mathbf{A}_*^{-1/2}\tilde{\mathbf{u}} \right\rangle_{\mathcal{H}^n} = \text{tr} \left(\tilde{\mathbf{P}}\mathbf{A}_*^{-1} \right) = \text{tr}(\mathbf{A}_*\mathbf{P}), \\ \langle \tilde{\mathbf{u}}_*, \mathbf{u} \rangle_{\mathcal{H}^n} &= \langle \mathbf{A}_*\mathbf{u}, \mathbf{u} \rangle_{\mathcal{H}^n} = \text{tr}(\mathbf{A}_*\mathbf{P}), \end{aligned}$$

where the last equal sign in the second line is due to the second-moment matching constraint

$\mathbf{A}_* \mathbf{P} \mathbf{A}_* = \tilde{\mathbf{P}}$. Combining the above results, we have shown

$$\langle \tilde{\mathbf{u}}, \mathbf{u} \rangle_{\mathcal{H}^n} \leq \text{tr}(\mathbf{A}_* \mathbf{P}) = \langle \tilde{\mathbf{u}}_*, \mathbf{u} \rangle_{\mathcal{H}^n}$$

for any $\tilde{\mathbf{u}} \in \mathcal{S}_{\tilde{\mathbf{P}}}$, which completes the proof. \square

Note that for the proof's own sake, actually the second part itself will suffice, which finds a lower bound of the objective function and shows that it is attainable. This also applies to theorem 2.1. However, such a proof is not constructive and is predicated on knowing (2.12) a priori. If we want to proceed with the Cauchy-Schwarz inequality approach without knowing (2.12), the problem reduces to looking for linear transforms $\tilde{\mathbf{u}} = \tilde{\mathbf{B}}\tilde{\mathbf{v}}$ and $\mathbf{u} = \mathbf{B}\mathbf{v}$ that yield the maximum attainable lower bound for $\langle \tilde{\mathbf{u}}_*, \mathbf{u} \rangle_{\mathcal{H}^n}$ and showing that it is indeed the upper bound for all candidate $\langle \tilde{\mathbf{u}}, \mathbf{u} \rangle_{\mathcal{H}^n}$. However, this is not as natural as the direct optimization approach employed in the above, by Lagrangian multipliers or calculus of variation.

2.5 Connections to Optimal Transport

Our main results theorem 2.1 and theorem 2.6 are also closely connected to the problem of optimal transport (a.k.a. the Monge-Kantorovich minimization problem (Villani, 2008, p. 10)), which concerns the coupling (in the form of a joint distribution) between two given probability distributions that minimizes an associated transport cost. More precisely, in a particular setting that is relevant to our context, given two PDFs $\mu(\cdot)$ and $\nu(\cdot)$ with finite second moments over \mathbb{R}^n , the problem concerns the joint distribution $\pi(\cdot, \cdot)$ over $\mathbb{R}^n \times \mathbb{R}^n$, which can be degenerate, such that π has μ and ν as marginals and the L_2 distance between the two marginal random variables are minimized, i.e.

$$\pi_* = \arg \min_{\pi \in \Pi_{\mu, \nu}} \int_{\mathbb{R}^n \times \mathbb{R}^n} \|\mathbf{x} - \mathbf{y}\|^2 \pi(\mathbf{x}, \mathbf{y}) d\mathbf{x} d\mathbf{y}, \quad (2.30)$$

where $\|\cdot\|$ is the Euclidean distance in \mathbb{R}^n and

$$\Pi_{\mu, \nu} = \left\{ \pi(\cdot, \cdot) \in \text{PDFs} : \int_{\mathbb{R}^n} \pi(\mathbf{x}, \mathbf{y}) d\mathbf{y} = \mu(\mathbf{x}), \int_{\mathbb{R}^n} \pi(\mathbf{x}, \mathbf{y}) d\mathbf{x} = \nu(\mathbf{y}) \right\}. \quad (2.31)$$

Under the above conditions, a unique optimizer π_* exists (Villani, 2008, p. 11) and it turns out to be deterministic, i.e. $\pi_*(\mathbf{x}, \mathbf{y}) = \delta_{\mathbf{y}=T_*(\mathbf{x})}\mu(\mathbf{x})$ for some mapping $T_* : \mathbb{R}^n \rightarrow \mathbb{R}^n$ such that $\int \pi_*(\mathbf{x}, \mathbf{y})d\mathbf{x} = \nu(\mathbf{y})$. Moreover, the quantity

$$W_2(\mu, \nu) = \min_{\pi \in \Pi_{\mu, \nu}} \left(\int_{\mathbb{R}^n \times \mathbb{R}^n} \|\mathbf{x} - \mathbf{y}\|^2 \pi(\mathbf{x}, \mathbf{y}) d\mathbf{x} d\mathbf{y} \right)^{1/2} \quad (2.32)$$

is called the L_2 Wasserstein distance between μ and ν (Villani, 2008, ch. 6).

In general, the optimal transport mapping $T_*(\cdot)$ is nonlinear. However, if we relax the constraint of marginal matching to only first- and second-moment matching, theorem 2.6 implies that $T_*(\cdot)$ is indeed affine with the linear part being SPD, even when μ and ν do not have a well-defined PDF (i.e. not absolutely continuous with respect to the Lebesgue measure). To show this, first notice that the optimization (2.30) over joint distributions on $\mathbb{R}^n \times \mathbb{R}^n$ can be translated into an equivalent problem over random vectors in \mathbb{R}^n :

$$\tilde{\mathbf{u}}_* = \arg \min_{\tilde{\mathbf{u}} \sim \nu} \mathbb{E} [\|\tilde{\mathbf{u}} - \mathbf{u}\|^2] \quad (2.33)$$

with some $\mathbf{u} \sim \mu$. Next, denote the mean of μ and ν by \mathbf{m}_μ and \mathbf{m}_ν , respectively, and also the covariance matrices by Σ_μ and Σ_ν . Then the objective function in (2.33) becomes

$$\mathbb{E} [\|\tilde{\mathbf{u}} - \mathbf{u}\|^2] = \|\mathbf{m}_\mu - \mathbf{m}_\nu\|^2 + \mathbb{E} [\|\tilde{\mathbf{u}}' - \mathbf{u}'\|^2],$$

where the primes indicate the centered random vectors. Now if we relax the distribution matching constraint $\tilde{\mathbf{u}} \sim \nu$ to only mean and covariance matching, i.e. $\mathbb{E}[\tilde{\mathbf{u}}] = \mathbf{m}_\nu$ and $\mathbb{E}[\tilde{\mathbf{u}}'\tilde{\mathbf{u}}'^T] = \Sigma_\nu$, the optimal transport problem (2.33) is simplified to the task of minimum-correction second-moment matching:

$$\tilde{\mathbf{u}}'_* = \arg \min_{\mathbb{E}[\tilde{\mathbf{u}}'\tilde{\mathbf{u}}'^T] = \Sigma_\nu} \mathbb{E} [\|\tilde{\mathbf{u}} - \mathbf{u}\|^2]. \quad (2.34)$$

If we set $\mathcal{H} = L_2(\Omega)$ as the Hilbert space of all square-integrable random variables defined over the probability space $(\Omega, \mathcal{A}, \mathcal{P})$ on which \mathbf{u} is defined, then theorem 2.6 applies to (2.34) and yields $\tilde{\mathbf{u}}_* = T_*(\mathbf{u}) = \mathbf{A}_*(\mathbf{u} - \mathbf{m}_\mu) + \mathbf{m}_\nu$ with

$$\mathbf{A}_* = \sqrt{\Sigma_\mu}^{-1} (\sqrt{\Sigma_\mu} \Sigma_\nu \sqrt{\Sigma_\mu}^T)^{1/2} \sqrt{\Sigma_\mu}^{-T}. \quad (2.35)$$

Moreover, since

$$\mathbb{E} [\|\tilde{\mathbf{u}}' - \mathbf{u}'\|^2] = \text{tr}(\Sigma_\mu) + \text{tr}(\Sigma_\nu) - 2\text{tr}[\text{Cov}(\tilde{\mathbf{u}}, \mathbf{u})] \quad (2.36)$$

and

$$\text{tr}[\text{Cov}(\tilde{\mathbf{u}}_*, \mathbf{u})] = \text{tr}[\text{Cov}(\mathbf{A}_* \mathbf{u}', \mathbf{u}')] = \text{tr}[(\sqrt{\Sigma_\mu} \Sigma_\nu \sqrt{\Sigma_\mu}^T)^{1/2}] = \text{tr}[(\Sigma_\mu \Sigma_\nu)^{1/2}],$$

we have

$$W_2^2(\mu, \mathbf{T}_* \mu) = \|\mathbf{m}_\mu - \mathbf{m}_\nu\|^2 + \text{tr}[\Sigma_\mu + \Sigma_\nu - 2(\Sigma_\mu \Sigma_\nu)^{1/2}], \quad (2.37)$$

where the distribution $\mathbf{T}_* \mu$ is the push-forward of μ by \mathbf{T}_* . Note that although $\mathbf{T}_* \mu$ share the same mean and covariance with ν , in general $\mathbf{T}_* \mu \neq \nu$ and $W_2(\mu, \mathbf{T}_* \mu) < W_2(\mu, \nu)$, so (2.37) indeed provides a lower bound for $W_2(\mu, \nu)$ based on the first two moments of μ and ν .

If we partition the set of all distributions (possibly degenerate) over \mathbb{R}^n that have finite second moments into equivalent classes by their means and covariances then (2.37) induces a metric on this quotient space. Denote by $\mathcal{S}_{\mathbf{m}, \Sigma}$ the equivalent class of all distributions with mean \mathbf{m} and covariance Σ . Given $\mathcal{S}_{\mathbf{m}_2, \Sigma_2}$ and $\mu \in \mathcal{S}_{\mathbf{m}_1, \Sigma_1}$, there exists a unique $\tilde{\mu} = \mathbf{T}_* \mu \in \mathcal{S}_{\mathbf{m}_2, \Sigma_2}$ that is the closest to μ under the distance $W_2(\cdot, \cdot)$. As a generalization to the one-to-one correspondence of (2.14) and the intuition illustrated in fig. 2-1, the members in any two equivalent classes $\mathcal{S}_{\mathbf{m}_1, \Sigma_1}$ and $\mathcal{S}_{\mathbf{m}_2, \Sigma_2}$ can be paired up by such minimum- W_2 -distance matching. Moreover, since this distance between μ and $\mathbf{T}_* \mu$ depends only on $\mathcal{S}_{\mathbf{m}_1, \Sigma_1}$ and $\mathcal{S}_{\mathbf{m}_2, \Sigma_2}$ but not on the particular member μ , $W_2(\mu, \mathbf{T}_* \mu)$ can be used to define the distance between $\mathcal{S}_{\mathbf{m}_1, \Sigma_1}$ and $\mathcal{S}_{\mathbf{m}_2, \Sigma_2}$. This is sometimes called the Fréchet distance (Dowson and Landau, 1982). Note that the covariance part of (2.37) can also be isolated and used as a metric among SPD matrices.

Although in general (2.37) is only a lower bound for $W_2(\mu, \nu)$, if both μ and ν belong to a family of distributions closed under affine transforms and each member of this family is uniquely determined by its mean and covariance, then we have $\mathbf{T}_* \mu = \nu$ and thus $W_2(\mu, \nu)$ coincides with (2.37). Dowson and Landau (1982) mentions one typical example of such where the distribution family is characterized by a radial kernel $\phi : [0, \infty) \rightarrow [0, \infty)$ that satisfies $0 < \int_0^\infty r^{2n+1} \phi(r^2) dr < \infty$, which guarantees the finiteness of second moments. In this family, a member μ 's PDF takes the form $\mu(\mathbf{x}) \propto \phi((\mathbf{x} - \mathbf{m})^T \mathbf{A} (\mathbf{x} - \mathbf{m}))$ with some SPD \mathbf{A} , where \mathbf{m} is the mean and \mathbf{A}^{-1} is a multiple of the covariance. In other words, this family is generated by pushing forward a non-degenerate elliptically contoured distribution that has finite second moments by invertible affine

transforms. It can be checked that within such a family, a distribution can be uniquely identified by its mean and covariance. Therefore, given μ and ν , we always have $T_*\mu = \nu$ and thus $W_2(\mu, \nu)$ can be computed by (2.37). The most familiar example of such is probably the family of Gaussian distributions corresponding to the exponential kernel $\phi(r) = e^{-r}$, for which the above results are well known.

To the best of our knowledge, the earliest work that derives the right hand side expression in (2.37) is in Olkin and Pukelsheim (1982) and Dowson and Landau (1982). However, they did not connect this result to the problem of optimal transport. Moreover, they translate the optimization into one over the cross-covariance matrix $B = \text{Cov}(\mathbf{u}, \tilde{\mathbf{u}})$, which does not necessarily uniquely determine $\tilde{\mathbf{u}}$ given \mathbf{u} if the optimizer B_* does not correspond to a deterministic affine coupling between the two. In theorem 2.6, although the objective function is also simplified to $\text{tr}(\text{Cov}(\mathbf{u}, \tilde{\mathbf{u}}))$ (or its empirical counterpart (2.1) in the discrete case of theorem 2.1), we have provided an alternative proof based on optimizing on $\tilde{\mathbf{u}}$ directly rather than on $B = \text{Cov}(\mathbf{u}, \tilde{\mathbf{u}})$.

Finally, note also that the more general result theorem 2.6 on a Hilbert space indeed contains the special case on \mathbb{R}^m described by theorem 2.1. Moreover, if we view theorem 2.1 through the lens of theorem 2.6, we can identify a discrete counterpart of the optimal transport interpretation for theorem 2.1. The key observation is that theorem 2.1 is equivalent to theorem 2.6 with $\mathcal{H} = \mathbb{R}^m$, which on the other hand, corresponds to (2.34) with μ being a discrete distribution over \mathbb{R}^n with m particles, i.e. with its (generalized) PDF being the sum of m Dirac deltas. Therefore, the optimal transport interpretation of theorem 2.6 implies that $\tilde{\mathbf{u}}_*$ also follows a discrete distribution with its m particles given by the rows of \tilde{U}_* in theorem 2.1. This is actually a stronger result than theorem 2.1 because in the latter, we have restricted the feasible $\tilde{\mu}$'s within the m -particle discrete distributions a priori.

For an exposition of other connections between the optimal transport and the polar decomposition, see Brenier (1991); Benamou and Brenier (2000).

2.6 Complications due to Degeneracy

In section 2.1, we have restricted ourselves to the nice full-rank case with $m \geq n = \text{rank}(\tilde{\mathbf{P}}) = \text{rank}(\mathbf{U})$. In practice, sometimes this is not satisfied, such as in some scenarios of data assimilation where the number m of samples of a random vector is smaller than the dimension n of the random vector itself, as will be discussed in chapter 6. Hence, here we will reveal the implications of degeneracy due to any of these full-rank assumptions being violated. Before moving on, we introduce the notations $r = \text{rank}(\mathbf{U})$ and $\tilde{r} = \text{rank}(\tilde{\mathbf{P}})$.

Step 1 First of all, since $\tilde{r} \leq n$, when $m < n$, it is possible that $m < \tilde{r}$, in which case $\text{rank}(\tilde{\mathbf{U}}^T \tilde{\mathbf{U}}) \leq m < \tilde{r}$, so there exists no $\tilde{\mathbf{U}} \in \mathcal{M}_{m \times n}$ whose second-moment matrix can match $\tilde{\mathbf{P}}$ due to rank deficiency. If $m < \tilde{r}$, an easiest fix is to find a best rank- m proxy to $\tilde{\mathbf{P}}$. For example, the leading rank- m truncation of the SVD of $\tilde{\mathbf{P}}$ (which is also the EVD here since $\tilde{\mathbf{P}}$ is semi-SPD) serves as the closest rank- m approximation under any unitarily invariant norm Higham (1988). See (Higham, 1988, sec. 6) for a review on other variants of a nearest lower-rank approximation. With $\tilde{\mathbf{P}}$ replaced by such a proxy when $m < \tilde{r}$, we can always ensure $m \geq \tilde{r}$ for the modified problem.

Step 2 Next, given $m \geq \tilde{r}$, if $\tilde{r} < n$, a variable substitution by a projection from \mathbb{R}^n onto $\text{Row}(\tilde{\mathbf{P}})$ will eliminate the rank deficiency in $\tilde{\mathbf{P}}$. If we construct $\tilde{\mathbf{Z}}$ with its columns forming an orthonormal basis of $\text{Row}(\tilde{\mathbf{P}})$, we have $\tilde{\mathbf{P}} = \tilde{\mathbf{Z}} \tilde{\mathbf{P}}_{\tilde{r}} \tilde{\mathbf{Z}}^T$ such that $\tilde{\mathbf{P}}_{\tilde{r}} = \tilde{\mathbf{Z}}^T \tilde{\mathbf{P}} \tilde{\mathbf{Z}} \in \mathcal{M}_{\tilde{r} \times \tilde{r}}$ has full rank. To have $\tilde{\mathbf{U}}^T \tilde{\mathbf{U}} = \tilde{\mathbf{P}}$, we must require $\text{Row}(\tilde{\mathbf{U}}) = \text{Row}(\tilde{\mathbf{P}})$, so $\tilde{\mathbf{U}}$ can be uniquely represented by $\tilde{\mathbf{S}} \in \mathcal{M}_{m \times \tilde{r}}$ such that $\tilde{\mathbf{U}} = \tilde{\mathbf{S}} \tilde{\mathbf{Z}}^T$. If we complete $\tilde{\mathbf{Z}}$ into an orthogonal $[\tilde{\mathbf{Z}}, \tilde{\mathbf{Z}}_{\perp}] \in \mathcal{M}_{n \times n}$, since

$$\|\tilde{\mathbf{U}} - \mathbf{U}\|_{\text{F}}^2 = \|(\tilde{\mathbf{U}} - \mathbf{U})[\tilde{\mathbf{Z}}, \tilde{\mathbf{Z}}_{\perp}]\|_{\text{F}}^2 = \|\tilde{\mathbf{S}} - \mathbf{U} \tilde{\mathbf{Z}}\|_{\text{F}}^2 + \|\mathbf{U} \tilde{\mathbf{Z}}_{\perp}\|_{\text{F}}^2,$$

the original optimization problem (1.4) can be equivalently formulated as

$$\arg \min_{\tilde{\mathbf{S}} \in \mathcal{M}_{m \times \tilde{r}}: \tilde{\mathbf{S}}^T \tilde{\mathbf{S}} = \tilde{\mathbf{Z}}^T \tilde{\mathbf{P}} \tilde{\mathbf{Z}}} \|\tilde{\mathbf{S}} - \mathbf{S}\|_{\text{F}}^2 \quad (2.38)$$

with $\mathbf{S} = \mathbf{U} \tilde{\mathbf{Z}}$. Since we already have $m \geq \tilde{r}$, with this further manipulation, we can always reduce a degenerate problem into one with $m \geq n = \text{rank}(\tilde{\mathbf{P}})$.

Step 3 Finally, given $m \geq n = \text{rank}(\tilde{P})$, let's consider the case of $\text{rank}(U) = r < n$. If we construct Z with its columns forming an orthonormal basis of $\text{Row}(U)$, U can be uniquely represented by a full-column-rank $W \in \mathcal{M}_{m \times r}$ such that $U = WZ^T$. Moreover, since Z can be augmented to an orthogonal $[Z, Z_\perp] \in \mathcal{M}_{n \times n}$, \tilde{U} can be uniquely represented by $\tilde{W} = \tilde{U}Z$ and $\tilde{W}_\perp = \tilde{U}Z_\perp$, which yield $\tilde{U} = \tilde{W}Z^T + \tilde{W}_\perp Z_\perp^T$ and

$$\|\tilde{U} - U\|_F^2 = \|(\tilde{U} - U)[Z, Z_\perp]\|_F^2 = \|\tilde{W} - W\|_F^2 + \|\tilde{W}_\perp\|_F^2.$$

Note that $\|\tilde{W}_\perp\|_F^2 = \text{tr}(Z_\perp^T \tilde{P} Z_\perp)$ is a constant thanks to the second-moment constraint. Therefore, the original problem (1.4) can be reduced to

$$\arg \min_{\tilde{W}, \tilde{W}_\perp: \tilde{U}^T \tilde{U} = \tilde{P}} \|\tilde{W} - W\|_F^2, \quad (2.39)$$

where $\tilde{U}^T \tilde{U} = \tilde{P}$ if and only if

$$\tilde{W}^T \tilde{W} = Z^T \tilde{P} Z, \quad \tilde{W}^T \tilde{W}_\perp = Z^T \tilde{P} Z_\perp, \quad \tilde{W}_\perp^T \tilde{W}_\perp = Z_\perp^T \tilde{P} Z_\perp. \quad (2.40)$$

Note the subtlety that although the objective function depends on \tilde{W} only, it is not obvious whether we can decouple the optimization over \tilde{W} and that over \tilde{W}_\perp , because it is possible that for the optimizer \tilde{W}_* to

$$\arg \min_{\tilde{W} \in \mathcal{M}_{m \times r}: \tilde{W}^T \tilde{W} = Z^T \tilde{P} Z} \|\tilde{W} - W\|_F^2, \quad (2.41)$$

there exists no \tilde{W}_\perp such that the second and third constraint in (2.40) are satisfied. These two matrix constraints consist of $r(n-r)$ linear scalar equations and $(n-r)(n-r+1)/2$ quadratic ones, which amount to $(n-r)(n-r+1)/2$ independent scalar constraints. On the other hand, we have $(n-r)m$ degrees of freedom in \tilde{W}_\perp . Therefore, under our assumption of $m \geq n \geq (r+1)$, we have $(n-r)(n-r+1)/2 \leq (n-r)n \leq (n-r)m$ so there should be enough degrees of freedom to ensure the existence of a valid \tilde{W}_\perp if all the constraints are compatible with each other. In the following, we will show that indeed for any \tilde{W} satisfying the first constraint in (2.40), there exists a valid \tilde{W}_\perp satisfying the other two. Moreover, we will characterize the set of all feasible \tilde{W}_\perp 's. This justifies reducing (2.39) to (2.41), which is of exactly the same form as (1.4) and does satisfy all the full-rank assumptions we have made for (1.4).

To identify all the solutions satisfying the last two constraints in (2.40), first we introduce the notation

$$\begin{bmatrix} B_{11} & B_{12} \\ B_{12}^T & B_{22} \end{bmatrix} \triangleq \begin{bmatrix} \tilde{W}^T \\ \tilde{W}_\perp^T \end{bmatrix} \begin{bmatrix} \tilde{W} & \tilde{W}_\perp \end{bmatrix} = \begin{bmatrix} Z^T \\ Z_\perp^T \end{bmatrix} \tilde{P} \begin{bmatrix} Z & Z_\perp \end{bmatrix}.$$

The general solution to the linear constraint $\tilde{W}^T \tilde{W}_\perp = B_{12}$ can be written as $\tilde{W}_\perp = (\tilde{W}^T)^+ B_{12} + \tilde{W}_c V$ with “+” denoting the Moore-Penrose pseudo-inverse (a.k.a. the generalized inverse), the columns of $\tilde{W}_c \in \mathcal{M}_{m \times (m-r)}$ forming an orthonormal basis of $\text{Col}(\tilde{W})^\perp$, and $V \in \mathcal{M}_{(m-r) \times (n-r)}$ arbitrary. Since B_{11} has full rank r , \tilde{W} has full column rank r and thus $\tilde{W}^+ = (\tilde{W}^T \tilde{W})^{-1} \tilde{W}^T = B_{11}^{-1} \tilde{W}^T$ and $(\tilde{W}^T)^+ = (\tilde{W}^+)^T = \tilde{W} B_{11}^{-1}$. Hence, $\tilde{W}_\perp = \tilde{W} B_{11}^{-1} B_{12} + \tilde{W}_c V$. Now the quadratic constraint $\tilde{W}_\perp^T \tilde{W}_\perp = B_{22}$ is equivalent to $V^T V = B_{22} - B_{12}^T B_{11}^{-1} B_{12}$, where the right hand side is an SPD Schur complement. A valid V exists if and only if $(m-r) \geq (n-r) \geq 1$. Therefore, the above analysis confirms our previous speculation based on comparing the number of degrees of freedom with the number of constraints. Besides, when a valid V exists, it is not unique.

We summarize the above three steps in the following theorem.

Theorem 2.7 (Degenerate counterpart of theorem 2.1). *Given $U \in \mathcal{M}_{m \times n}$ with $\text{rank}(U) = r$ and $\tilde{P} \in \mathcal{M}_{n \times n}$ with $\text{rank}(\tilde{P}) = \tilde{r} \leq m$, we have*

$$\arg \min_{\tilde{U} \in \mathcal{M}_{m \times n}: \tilde{U}^T \tilde{U} = \tilde{P}} \|\tilde{U} - U\|_F^2 = \tilde{U}_* = \tilde{S}_* \tilde{Z}^T = (\tilde{W}_* Z^T + \tilde{W}_{\perp*} Z_\perp^T) \tilde{Z}^T, \quad (2.42)$$

where $\tilde{S}_* = \tilde{W}_* Z^T + \tilde{W}_{\perp*} Z_\perp^T \in \mathcal{M}_{m \times \tilde{r}}$. Moreover, the columns of $\tilde{Z} \in \mathcal{M}_{n \times \tilde{r}}$ form an orthonormal basis of $\text{Row}(\tilde{P})$ and the columns of $Z \in \mathcal{M}_{\tilde{r} \times r'}$ form an orthonormal basis of $\text{Row}(U \tilde{Z})$ with

$$r' = \tilde{r} - \dim(\text{Row}(\tilde{P}) \cap \text{Row}(U)^\perp) = r - \dim(\text{Row}(U) \cap \text{Row}(\tilde{P})^\perp). \quad (2.43)$$

Finally, \tilde{W}_* and $\tilde{W}_{\perp*}$ are given by

$$\tilde{W}_* = \arg \min_{\tilde{W} \in \mathcal{M}_{m \times r'}: \tilde{W}^T \tilde{W} = Z^T \tilde{Z}^T \tilde{P} \tilde{Z} Z} \|\tilde{W} - U \tilde{Z} Z\|_F^2 \quad (2.44)$$

and $\tilde{W}_{\perp*} = \tilde{W}_* B_{11}^{-1} B_{12} + \tilde{W}_{c*} V_* \in \mathcal{M}_{m \times (\tilde{r}-r')}$ with the columns of $\tilde{W}_{c*} \in \mathcal{M}_{m \times (m-r')}$ forming an orthonormal basis of $\text{Col}(\tilde{W}_*)^\perp$ and $V_* \in \mathcal{M}_{(m-r') \times (\tilde{r}-r')}$ arbitrary as long as $V_*^T V_* = B_{22} -$

$B_{12}^T B_{11}^{-1} B_{12}$. The B_{ij} blocks are defined as:

$$\begin{bmatrix} B_{11} & B_{12} \\ B_{12}^T & B_{22} \end{bmatrix} \triangleq \begin{bmatrix} \tilde{W}^T \\ \tilde{W}_\perp^T \end{bmatrix} \begin{bmatrix} \tilde{W} & \tilde{W}_\perp \end{bmatrix} = \begin{bmatrix} Z^T \\ Z_\perp^T \end{bmatrix} \tilde{Z}^T \tilde{P} \tilde{Z} \begin{bmatrix} Z & Z_\perp \end{bmatrix}.$$

Besides, the global minimum can be expressed in exactly the same way as in (2.11):

$$\|\tilde{U}_* - U\|_F^2 = \text{tr}[P + \tilde{P} - 2(P\tilde{P})^{1/2}]. \quad (2.45)$$

Proof. Everything other than (2.45) ensues readily from knitting into one piece step 2 and step 3 shown right before we state this theorem. To prove (2.45), notice that

$$\|\tilde{U}_* - U\|_F^2 - \text{tr}(P) - \text{tr}(\tilde{P}) = -2\text{tr}(U^T \tilde{U}_*) = -2\text{tr}((U\tilde{Z})^T \tilde{S}_*) = -2\text{tr}((U\tilde{Z}Z)^T \tilde{W}_*).$$

Since theorem 2.1 applies to (2.44), the last expression in the above can be computed by (2.11) to yield $\|\tilde{U}_* - U\|_F^2 = \text{tr}(P) + \text{tr}(\tilde{P}) - 2\text{tr}[(P_{r'}\tilde{P}_{r'})^{1/2}]$, where $P_{r'} = Z^T \tilde{Z}^T P \tilde{Z} Z$ and $\tilde{P}_{r'} = Z^T \tilde{Z}^T \tilde{P} \tilde{Z} Z$. Hence, what remains to be shown is $\text{tr}[(P_{r'}\tilde{P}_{r'})^{1/2}] = \text{tr}[(P\tilde{P})^{1/2}]$ and it suffices to show that $P_{r'}\tilde{P}_{r'}$ and $P\tilde{P}$ share the same nonzero eigenvalues with the same algebraic multiplicity. To demonstrate this, first compute

$$P_{r'}\tilde{P}_{r'} = Z^T \tilde{Z}^T P \tilde{Z} (Z Z^T \tilde{Z}^T \tilde{P}) \tilde{Z} Z = Z^T \tilde{Z}^T P (\tilde{Z} \tilde{Z}^T \tilde{P}) \tilde{Z} Z = Z^T \tilde{Z}^T P \tilde{P} \tilde{Z} Z$$

using the property of the orthogonal projections ZZ^T and $\tilde{Z}\tilde{Z}^T$. Next, notice that $(Z^T \tilde{Z}^T P)(\tilde{P} \tilde{Z} Z)$ and $(\tilde{P} \tilde{Z} Z)(Z^T \tilde{Z}^T P)$ have identical nonzero eigenvalues with the same algebraic multiplicity. Since $(\tilde{P} \tilde{Z} Z)(Z^T \tilde{Z}^T P) = \tilde{P} \tilde{Z} \tilde{Z}^T P = \tilde{P} P$, we conclude that the nonzero spectrum of $P_{r'}\tilde{P}_{r'}$ and of $P\tilde{P}$ coincide in terms of eigenvalues and their algebraic multiplicity, which completes the proof. \square

Remark 1 \tilde{U}_* is unique if and only if $r' = \tilde{r}$, i.e. $\text{Row}(\tilde{P}) \cap \text{Row}(U)^\perp = 0$. In this case, step 3 is not needed. Given $r' = \tilde{r}$, the reverse problem seeking the closest point to \tilde{U}_* with second-moment matrix P will also have a unique solution if and only if $r' = r$, i.e. $\text{Row}(U) \cap \text{Row}(\tilde{P})^\perp = 0$. Therefore, the elements in $\mathcal{S}_P = \{U \in \mathcal{M}_{m \times n} : U^T U = P\}$ and those in $\mathcal{S}_{\tilde{P}}$ can be paired up bijectively and unambiguously by this closest-point correspondence if and only if $\text{rank}(P) = \text{rank}(\tilde{P}) = \text{rank}(P\tilde{P})$. In practice, one common special case of such is when $\text{Row}(U) = \text{Row}(\tilde{P})$,

as will be seen in chapter 6. If this is true, besides fixing rank deficiency, step 2 also preserves the value of the objective function, i.e. $\|\tilde{\mathbf{S}} - \mathbf{S}\|_{\mathbb{F}}^2 = \|\tilde{\mathbf{U}}\tilde{\mathbf{Z}} - \mathbf{U}\tilde{\mathbf{Z}}\|_{\mathbb{F}}^2 \equiv \|\tilde{\mathbf{U}} - \mathbf{U}\|_{\mathbb{F}}^2$.

Remark 2 One extreme case of rank deficiency is when $\text{Row}(\tilde{\mathbf{P}}) \perp \text{Row}(\mathbf{U})$ and thus $r' = 0$. This renders the optimization (2.42) trivial because the objective function is now constant: $\|\tilde{\mathbf{U}} - \mathbf{U}\|_{\mathbb{F}}^2 \equiv \text{tr}(\mathbf{P} + \tilde{\mathbf{P}})$, so all feasible $\tilde{\mathbf{U}}$'s are equally good.

Remark 3 Although the degeneracy significantly complicates the characterization of the optimizer $\tilde{\mathbf{U}}_*$, the minimum of the objective function (2.45) surprisingly shares the same simple formula as (2.11) in the non-degenerate case.

theorem 2.7 can also be generalized to the case with \mathbb{R}^m replaced by a Hilbert space \mathcal{H} . This continuous setup can be intuitively viewed as the discrete one with $m = \infty$, which implies that we always have $m > n \geq \tilde{r}$ so the modification in step 1 is never needed. We state this generalization in the following theorem, whose proof is omitted since it is totally in parallel with that of theorem 2.7.

Theorem 2.8 (Degenerate counterpart of theorem 2.6). *Given $\mathbf{u} \in \mathcal{H}^n$ with $\mathbf{P} = \text{Gr}(\mathbf{u}, \mathbf{u}) \in \mathcal{M}_{n \times n}$ and $\text{rank}(\mathbf{P}) = \dim(\text{span}\{\mathbf{u}_1, \dots, \mathbf{u}_n\}) = r$, and semi-SPD $\tilde{\mathbf{P}} \in \mathcal{M}_{n \times n}$ with $\text{rank}(\tilde{\mathbf{P}}) = \tilde{r}$, we have*

$$\arg \min_{\tilde{\mathbf{u}} \in \mathcal{H}^n: \text{Gr}(\tilde{\mathbf{u}}, \tilde{\mathbf{u}}) = \tilde{\mathbf{P}}} \|\tilde{\mathbf{u}} - \mathbf{u}\|_{\mathcal{H}^n}^2 = \tilde{\mathbf{u}}_* = \tilde{\mathbf{Z}}\tilde{\mathbf{s}}_* = \tilde{\mathbf{Z}}(\mathbf{Z}\tilde{\mathbf{w}}_* + \mathbf{Z}_{\perp}\tilde{\mathbf{w}}_{\perp*}) \quad (2.46)$$

where $\tilde{\mathbf{s}}_* = \mathbf{Z}\tilde{\mathbf{w}}_* + \mathbf{Z}_{\perp}\tilde{\mathbf{w}}_{\perp*} \in \mathcal{H}^{\tilde{r}}$. Moreover, the columns of $\tilde{\mathbf{Z}} \in \mathcal{M}_{n \times \tilde{r}}$ form an orthonormal basis of $\text{Row}(\tilde{\mathbf{P}})$ and the columns of $\mathbf{Z} \in \mathcal{M}_{\tilde{r} \times r'}$ form an orthonormal basis of $\text{Row}(\tilde{\mathbf{Z}}^{\text{T}}\mathbf{P}\tilde{\mathbf{Z}})$ with

$$r' = \tilde{r} - \dim(\text{Row}(\tilde{\mathbf{P}}) \cap \text{Row}(\mathbf{P})^{\perp}) = r - \dim(\text{Row}(\mathbf{P}) \cap \text{Row}(\tilde{\mathbf{P}})^{\perp}). \quad (2.47)$$

Finally, $\tilde{\mathbf{w}}_*$ and $\tilde{\mathbf{w}}_{\perp*}$ are given by

$$\tilde{\mathbf{w}}_* = \arg \min_{\tilde{\mathbf{w}} \in \mathcal{H}^{r'}: \text{Gr}(\tilde{\mathbf{w}}, \tilde{\mathbf{w}}) = \mathbf{Z}^{\text{T}}\tilde{\mathbf{Z}}^{\text{T}}\tilde{\mathbf{P}}\tilde{\mathbf{Z}}} \|\tilde{\mathbf{w}} - \mathbf{Z}^{\text{T}}\tilde{\mathbf{Z}}^{\text{T}}\mathbf{u}\|_{\mathcal{H}^{r'}}^2 \quad (2.48)$$

and $\tilde{\mathbf{w}}_{\perp*} = \mathbf{B}_{12}^T \mathbf{B}_{11}^{-1} \tilde{\mathbf{w}}_* + \mathbf{v}_* \in \mathcal{H}^{\tilde{r}-r'}$ with arbitrary \mathbf{v}_* such that $\text{Gr}(\mathbf{v}_*, \tilde{\mathbf{w}}_*) = \mathbf{0}$, i.e. $\mathbf{v}_{*,i} \in \text{span}\{\tilde{\mathbf{w}}_{*,1}, \dots, \tilde{\mathbf{w}}_{*,r'}\}^\perp \subset \mathcal{H}$ for $i = 1, \dots, (\tilde{r} - r')$, and $\text{Gr}(\mathbf{v}_*, \mathbf{v}_*) = \mathbf{B}_{22} - \mathbf{B}_{12}^T \mathbf{B}_{11}^{-1} \mathbf{B}_{12}$. The \mathbf{B}_{ij} blocks are defined as:

$$\begin{bmatrix} \mathbf{B}_{11} & \mathbf{B}_{12} \\ \mathbf{B}_{12}^T & \mathbf{B}_{22} \end{bmatrix} \triangleq \text{Gr} \left(\begin{bmatrix} \tilde{\mathbf{w}} \\ \tilde{\mathbf{w}}_{\perp} \end{bmatrix}, \begin{bmatrix} \tilde{\mathbf{w}} \\ \tilde{\mathbf{w}}_{\perp} \end{bmatrix} \right) = \begin{bmatrix} \mathbf{Z}^T \\ \mathbf{Z}_{\perp}^T \end{bmatrix} \tilde{\mathbf{Z}}^T \tilde{\mathbf{P}} \tilde{\mathbf{Z}} \begin{bmatrix} \mathbf{Z} \\ \mathbf{Z}_{\perp} \end{bmatrix}.$$

Besides, the global minimum of the objective function is

$$\|\tilde{\mathbf{u}}_* - \mathbf{u}\|_{\mathcal{H}^n}^2 = \text{tr}[\mathbf{P} + \tilde{\mathbf{P}} - 2(\mathbf{P}\tilde{\mathbf{P}})^{1/2}]. \quad (2.49)$$

All previous remarks regarding theorem 2.7 apply here as well. Moreover, this theorem also extends our optimal transport interpretation in section 2.5 to the most general case with arbitrary semi-SPD covariance matrix Σ_μ and Σ_ν , unlike most previous literature (e.g. Dowson and Landau (1982); Olkin and Pukelsheim (1982)) on this topic, which essentially restrict themselves in cases with $\text{Row}(\Sigma_\nu) \subset \text{Row}(\Sigma_\mu)$. In particular, theorem 2.8 implies that the optimal coupling π_* between μ and ν (which is a joint distribution with μ and ν being its marginals) for matching only the first two moments is deterministic (i.e. corresponds to a bijection between $\mathbf{u} \sim \mu$ and $\tilde{\mathbf{u}} \sim \nu$) if and only if $\text{rank}(\Sigma_\mu) = \text{rank}(\Sigma_\nu) = \text{rank}(\Sigma_\mu \Sigma_\nu)$. The optimal coupling is deterministic in the direction from μ to ν (i.e. π_* induces a map from \mathbf{u} to $\tilde{\mathbf{u}}$ or the conditional $\pi_{\nu|\mu}$ is always a Dirac delta) if and only if $\text{Row}(\Sigma_\nu) \cap \text{Row}(\Sigma_\mu)^\perp = 0$, of which the assumption $\text{Row}(\Sigma_\nu) \subset \text{Row}(\Sigma_\mu)$ usually made in previous literature is a special case.

2.7 Maximization Counterparts

Finally, we remark that although our motivation is to find the minimum correction for second-moment matching, all the above results indeed have their maximization counterparts, which are summarized in corollary 2.9. They can be proved in exactly the same way as their minimization counterparts. The geometric picture shown in fig. 2-1 also hints on why a global maximizer exists and is simply the negative of the minimizer.

Corollary 2.9 (Maximization counterparts). *With the same assumptions required in theorem 2.1 and theorem 2.6, we have*

$$\arg \max_{\tilde{U} \in \mathcal{M}_{m \times n}: \tilde{U}^T \tilde{U} = \tilde{P}} \|\tilde{U} - U\|_{\mathbb{F}}^2 = -U A_*. \quad (2.50)$$

$$\arg \max_{\tilde{\mathbf{u}} \in \mathcal{H}^n: \text{Gr}(\tilde{\mathbf{u}}, \tilde{\mathbf{u}}) = \tilde{P}} \|\tilde{\mathbf{u}} - \mathbf{u}\|_{\mathcal{H}^n} = -A_* \mathbf{u}, \quad (2.51)$$

with the SPD A_* given by (2.12).

Chapter 3

Algorithms for Minimum-Correction Second-Moment Matching

3.1 General Computation Strategies

To compute \tilde{U}_* in theorem 2.1, there are two routes. We can either compute $\tilde{U}_* = U A_*$ using

$$A_* = \sqrt{\tilde{P}}^T (\sqrt{\tilde{P}} P \sqrt{\tilde{P}}^T)^{-1/2} \sqrt{\tilde{P}} = \sqrt{P}^{-1} (\sqrt{P} \tilde{P} \sqrt{P}^T)^{1/2} \sqrt{P}^{-T} \quad (3.1)$$

or directly obtain \tilde{U}_* without explicitly forming P or computing A_* .

Computing \tilde{U}_* through A_* To take the A_* route, we need a pre-processing step $P = U^T U$ costing $2mn^2$ flops (floating point operations), a post-processing step $\tilde{U}_* = U A_*$ costing $2mn^2$ flops, and the key step (3.1).

In almost all previous results regarding optimal transport for Gaussian distributions (e.g. Dowson and Landau (1982); Olkin and Pukelsheim (1982); Ott et al. (2004) and (Reich and Cotter, 2015, p. 55)), the optimal mapping (3.1) is written in a form that every matrix square root is chosen to be the unique SPD one, i.e. $A_* = \tilde{P}^{1/2} (\tilde{P}^{1/2} P \tilde{P}^{1/2})^{-1/2} \tilde{P}^{1/2}$. However, our derivation of (3.1) indi-

cates that only one of these five square roots has to be the SPD one, while all the others can actually be replaced by an arbitrary square root. Analytically this is immaterial since \mathbf{A}_* is unchanged and we might as well simplify the expression by always using the SPD square root. However, numerically computing an arbitrary square root is much easier and cheaper than computing the SPD one. The algorithm of choice to compute an arbitrary square root is probably the Cholesky decomposition which gives an upper-triangular square root and costs $(1/3)n^3$ flops (Trefethen and Bau III, 1997, p. 176). This algorithm is known to be efficient, numerically stable, and widely available. Therefore, computing \mathbf{A}_* involves one Cholesky decomposition, one SPD square root computation (cost denoted by $C_{\text{sqr}}n^3$) and four matrix multiplications. The total cost of computing $\tilde{\mathbf{U}}_*$ is thus $4mn^2 + (8 + 1/3 + C_{\text{sqr}})n^3$.

Note that for modern vectorized and multi-threading computer architectures, the number of flops is no longer dominant in affecting the relative efficiency among algorithms sharing the same asymptotic cost growth rate. How well an algorithm can make use of time and space concurrency in cache and memory reference has become much more important. Therefore, here we count the flops simply for reference, but not to seriously compare efficiency. We will only count the leading order term(s) and we will omit symmetry. Moreover, we will not distinguish between the cost of \mathbf{AB} and of $\mathbf{A}^{-1}\mathbf{B}$ for square \mathbf{A} and \mathbf{B} when counting flops, although $\mathbf{A}^{-1}\mathbf{B}$ in general takes 50% more time than \mathbf{AB} (e.g. based on tests in MATLAB; see also Higham and Schreiber (1990)).

The drawback of the \mathbf{A}_* route is that if \mathbf{U} is ill-conditioned, forming $\mathbf{P} = \mathbf{U}^T\mathbf{U}$ indeed squares the condition number, i.e. $\kappa(\mathbf{P}) = \kappa(\mathbf{U})^2$, and renders the problem even much more ill-conditioned. Therefore, in an ill-conditioned case, the non- \mathbf{A}_* route is preferred.

Computing $\tilde{\mathbf{U}}_*$ without \mathbf{A}_* The non- \mathbf{A}_* route is based on computing the $\mathcal{S}_{\tilde{\mathbf{P}}}$ factor $\tilde{\mathbf{U}}_*$ of the generalized PD (polar decomposition) $\mathbf{U} = \tilde{\mathbf{U}}_*\mathbf{C}$ introduced in corollary 2.5. First, we need to reduce the rectangular PD to a square one by identifying an orthonormal basis of \mathbf{U} 's columns, i.e. $\mathbf{U} = \mathbf{V}\mathbf{R}$ with $\mathbf{V}^T\mathbf{V} = \mathbf{I}$ and $\mathbf{R} \in \mathcal{M}_{n \times n}$. This can be achieved, for example, by a Householder QR factorization costing $(2mn^2 - (2/3)n^3)$ flops (Trefethen and Bau III, 1997, ch. 10). Note that the classical or modified Gram-Schmidt will not yield an accurately orthogonal \mathbf{V} for an ill-conditioned \mathbf{U} (Trefethen and Bau III, 1997, ch. 9). Next we reduce computing the $\mathcal{S}_{\tilde{\mathbf{P}}}$ factor $\tilde{\mathbf{R}}_*$ of the generalized PD $\mathbf{R} = \tilde{\mathbf{R}}_*\mathbf{C}$ to computing the orthogonal polar factor \mathbf{Q} of the PD $\mathbf{R}\sqrt{\tilde{\mathbf{P}}}^T = \mathbf{Q}\mathbf{S} =$

$(\tilde{\mathbf{R}}_* \sqrt{\tilde{\mathbf{P}}}^{-1})(\sqrt{\tilde{\mathbf{P}}} \mathbf{C} \sqrt{\tilde{\mathbf{P}}}^T)$. Finally, we can assemble the desired $\tilde{\mathbf{U}}_*$ as $\tilde{\mathbf{U}}_* = \mathbf{V} \tilde{\mathbf{R}}_* = \mathbf{V} \mathbf{Q} \sqrt{\tilde{\mathbf{P}}}$. Therefore, if we denote by $C_{\text{po}} n^3$ the cost of computing the orthogonal polar factor of an n -by- n matrix, the total cost of obtaining $\tilde{\mathbf{U}}_*$ becomes $4mn^2 + (4 - (1/3) + C_{\text{po}})n^3$.

To summarize, the \mathbf{A}_* route is eventually reduced to computing an SPD square root, while the non- \mathbf{A}_* route is reduced to performing a polar decomposition. These two tasks are indeed intimately related and an algorithm for one task usually corresponds to one for the other, as will be seen in the next subsection.

In the above, we have assumed the problem is non-degenerate, as required by theorem 2.1. If that is not the case, the constructive proof of theorem 2.7 already provides an algorithm to reduce a degenerate case to a non-degenerate one.

3.2 Algorithms for the SPD Square Root and Polar Decomposition

Given an invertible $\mathbf{R} \in \mathcal{M}_{n \times n}$ with polar decomposition $\mathbf{R} = \mathbf{Q}\mathbf{C}$, where \mathbf{Q} and \mathbf{C} are the orthogonal and SPD polar factor, respectively, then \mathbf{C} is the SPD square root of $\mathbf{R}^T \mathbf{R}$, i.e. $\mathbf{C} = (\mathbf{R}^T \mathbf{R})^{1/2}$. On the other hand, given an SPD $\mathbf{P} \in \mathcal{M}_{n \times n}$ with an arbitrary square root $\sqrt{\mathbf{P}}$, $\mathbf{P}^{1/2}$ will also be the SPD polar factor of $\sqrt{\mathbf{P}}$. Therefore, computing an SPD square root and performing a polar decomposition can be reduced to each other in exactly the same vein as how the Cholesky decomposition is related to the QR factorization.

3.2.1 Algorithms Based on SVD/EVD

For the polar decomposition of $\mathbf{R} \in \mathcal{M}_{n \times n}$, a straightforward SVD-based algorithm stems from the fact that the SVD of \mathbf{R} can be readily manipulated to yield the polar decomposition:

$$\mathbf{R} = \mathbf{Q}\mathbf{\Sigma}\mathbf{V}^T = (\mathbf{Q}\mathbf{V}^T)(\mathbf{V}\mathbf{\Sigma}\mathbf{V}^T).$$

This procedure costs $(8/3)n^3$ flops for the SVD (Trefethen and Bau III, 1997, p. 237) and an extra $2n^3$ for forming either the orthogonal or the SPD polar factor. This SVD-based algorithm corresponds to using the EVD (eigenvalue decomposition) to obtain the SPD square root of an SPD $P \in \mathcal{M}_{n \times n}$ by $P = V\Lambda V^T = (V\Lambda^{1/2}V^T)^2$, which costs $(4/3)n^3$ flops for the EVD (Trefethen and Bau III, 1997, p. 200) and $2n^3$ for matrix multiplication.

Since we have efficient and numerically stable algorithms for SVD and EVD (Golub and Van Loan, 2013; Trefethen and Bau III, 1997), which are also widely available as highly optimized library functions in today's most computing platforms, the above SVD/EVD-based algorithms are easy to implement and robust to use. Despite this, there are cases where a non-SVD/EVD-based algorithm is desirable. It might be because, for example, we have no access to an efficient SVD/EVD procedure, we want to have more control over the accuracy of the result, or we want to utilize some special features particular to a problem at hand to further boost the efficiency. A non-SVD/EVD-based algorithm typically involves explicitly carrying out a matrix fixed-point iteration, as will be discussed next.

3.2.2 Algorithms Based on Fixed-Point Iterations

For computing the SPD square root and polar decomposition, there are also non-SVD/EVD-based algorithms that take the form of a matrix fixed-point iteration $X_{k+1} = F(X_k)$ with $F : \mathcal{M}_{n \times n} \rightarrow \mathcal{M}_{n \times n}$ involving only matrix addition, multiplication, and inversion.

If we view an iteration as a discrete dynamical system¹, the minimum requirements for a matrix iteration $X_{k+1} = F(X_k)$ to be usable in practice for seeking a target matrix X_* include:

1. X_* is a neutrally stable fixed point of F .
2. A subset \mathcal{A}_{ini} of the attraction basin $\mathcal{A}_{F, X_*} \subset \mathcal{M}_{n \times n}$ can be identified.

Here by definition, the attraction basin contains all the points that will converge to X_* under F , i.e. $X \in \mathcal{A}_{F, X_*}$ if and only if $X_k \rightarrow X_*$ with $X_0 = X$ and $X_{k+1} = F(X_k)$. If these two requirements are met, for any initial guess $X_0 \in \mathcal{A}_{\text{ini}}$, we have $X_k \rightarrow X_*$ analytically and the

¹See (Strogatz, 1994, sec. 5.1) for the definition of terms related to dynamical systems.

neutral stability of \mathbf{X}_* guarantees that any numerical errors incurred in carrying out $\mathbf{X}_{k+1} = \mathbf{F}(\mathbf{X}_k)$ will stay bounded once \mathbf{X}_k approaches \mathbf{X}_* , i.e. entering some ε -neighborhood of \mathbf{X}_* . However, technically there is still a minor gap because the numerical errors incurred in the early iterations may grow significantly before this ε -neighborhood of \mathbf{X}_* is reached and the neutral stability takes effect. Fortunately, this is usually not a problem in practice for iterations converging quadratically. If the convergence is only linear, the algorithm will become uncompetitive anyway simply due to its inefficiency.

As we shall see, usually we can find \mathcal{A}_{ini} as a stable manifold containing \mathbf{X}_* and invariant under \mathbf{F} . Sometimes \mathcal{A}_{ini} can even extend to infinity. However, this should not be confused with global convergence and it does not even guarantee the numerical stability of the algorithm because \mathcal{A}_{ini} as a submanifold of $\mathcal{M}_{n \times n}$ is a boundary set, i.e. it has no interior point. Hence, the numerical errors $(\hat{\mathbf{X}}_k - \mathbf{X}_k)$, although typically small in magnitude but in general unconstrained in structure, will drive the actual iterates $\hat{\mathbf{X}}_k$'s away from \mathcal{A}_{ini} and might make them fail to converge to anything close to \mathbf{X}_* . This is why we need the neutral stability of \mathbf{X}_* to ensure that if the iterates manage to approach \mathbf{X}_* before the numerical errors accumulate significantly, these errors will remain bounded hereafter. Note that a neutrally stable fixed point may not be attracting, so we don't even need \mathbf{X}_* to be an interior point of the attraction basin $\mathcal{A}_{\mathbf{F}, \mathbf{X}_*}$.

Since the dynamical system $\mathbf{X}_{k+1} = \mathbf{F}(\mathbf{X}_k)$ is n^2 -dimensional, we can reshape each \mathbf{X}_k into an $\mathbb{R}^{n \times n}$ vector and view \mathbf{F} as an operator from $\mathbb{R}^{n \times n}$ to itself. Denote the Jacobian matrix of \mathbf{F} at \mathbf{X}_* by $\mathbf{JF} \in \mathcal{M}_{n^2 \times n^2}^2$. Then a necessary condition for \mathbf{X}_* to be neutrally stable is $\rho(\mathbf{JF}) \leq 1$, where $\rho(\cdot)$ denotes the spectral radius, while the strict inequality $\rho(\mathbf{JF}) < 1$ is a sufficient one. However, viewing \mathbf{X}_k as a vector typically makes the expression of \mathbf{F} awkwardly complicated and renders cumbersome the analysis of \mathbf{JF} 's spectrum. Philippe (1987) bypasses this difficulty elegantly by keeping \mathbf{X}_k as a matrix and taking a functional viewpoint of $\mathbf{X}_{k+1} = \mathbf{F}(\mathbf{X}_k)$. From this perspective, the counterpart of \mathbf{JF} is the Fréchet derivative \mathbf{DF} of the operator \mathbf{F} at \mathbf{X}_* . Here $\mathbf{DF} : \mathcal{M}_{n \times n} \rightarrow \mathcal{M}_{n \times n}$ is a linear operator. Since \mathbf{JF} and \mathbf{DF} are essentially the same linear operator, they share the same spectrum. It turns out that identifying all the eigenvalues and eigenvectors of \mathbf{DF} by direct observation can be made possible by suitable linear transforms.

Next, we review the most common iterations for computing the i) SPD square root of an SPD

²The explicit dependence on \mathbf{X}_* is omitted in the notation since it should be clear from context.

matrix and ii) orthogonal polar factor of a square matrix. To do so, we analyze and compare iterations in a novel unified framework. We also obtain new insights into their properties and propose a slight modification that surprisingly remedies some iterations deemed numerically unstable previously.

i) A Unified Framework for Four SPD Square Root Iterations Given an SPD \mathbf{P} , the following four iterations can be used to compute the SPD square root $\mathbf{P}^{1/2}$ and $\mathbf{P}^{-1/2}$ (see Higham (1986b); Philippe (1987) and (Golub and Van Loan, 2013, sec. 9.4.2)):

$$\mathbf{X}_{k+1} = \mathbf{X}_k + \Delta \mathbf{X}_k, \quad \mathbf{X}_k \Delta \mathbf{X}_k + \Delta \mathbf{X}_k \mathbf{X}_k = \mathbf{P} - \mathbf{X}_k^2; \quad (3.2)$$

$$\mathbf{X}_{k+1} = (1/2)(\mathbf{X}_k + \mathbf{X}_k^{-1} \mathbf{P}); \quad (3.3)$$

$$\mathbf{X}_{k+1} = (1/2)(\mathbf{X}_k + \mathbf{Y}_k^{-1}), \quad \mathbf{Y}_{k+1} = (1/2)(\mathbf{Y}_k + \mathbf{X}_k^{-1}); \quad (3.4)$$

$$\mathbf{X}_{k+1} = \mathbf{X}_k + (1/2)\mathbf{X}_k(\mathbf{I} - \mathbf{X}_k \mathbf{P} \mathbf{X}_k). \quad (3.5)$$

Among these iterations, (3.2) stems from a direct application of the Newton's method to the matrix equation $\mathbf{G}(\mathbf{X}) = \mathbf{X}^2 - \mathbf{P} = \mathbf{0}$ with the Fréchet derivative of \mathbf{G} at \mathbf{X}_k being the linear operator $\mathbf{D}\mathbf{G}_{\mathbf{X}_k} \mathbf{X} = \mathbf{X}_k \mathbf{X} + \mathbf{X} \mathbf{X}_k$. (3.3) is a simplification of (3.2) by making the assumption $\mathbf{X}_k \Delta \mathbf{X}_k = \Delta \mathbf{X}_k \mathbf{X}_k$, which is true for certain choices of \mathbf{X}_0 as we shall see. (3.4) is an extension of (3.3) by coupling (3.3) with a mirror iteration to compute $\mathbf{P}^{1/2}$ and $\mathbf{P}^{-1/2}$ simultaneously. Finally, (3.5) is quite different from the others, as it comes from mimicking the Newton iteration for the scalar equation $g(x) = 1/x^2 - \lambda = 0$ (Philippe, 1987).

The converging point and a corresponding \mathcal{A}_{ini} of feasible initial values for iterations eqs. (3.2) to (3.5) are summarized as

$$(3.2), (3.3) : \mathbf{X}_k \rightarrow \mathbf{P}^{1/2} \quad \text{for } \mathbf{X}_0 \in \mathcal{C}_{\mathbf{P}} \triangleq \{\mathbf{X} \text{ SPD} : \mathbf{X}\mathbf{P} = \mathbf{P}\mathbf{X}\}; \quad (3.6)$$

$$(3.4) : \mathbf{X}_k \rightarrow \mathbf{P}^{1/2}, \quad \mathbf{Y}_k \rightarrow \mathbf{P}^{-1/2} \quad \text{for } \mathbf{X}_0 \in \mathcal{C}_{\mathbf{P}}, \quad \mathbf{Y}_0 = \mathbf{P}^{-1} \mathbf{X}_0; \quad (3.7)$$

$$(3.5) : \mathbf{X}_k \rightarrow \mathbf{P}^{-1/2} \quad \text{for } \mathbf{X}_0 \in \{\mathbf{X} \in \mathcal{C}_{\mathbf{P}} : \|\mathbf{X}\|_2 < \sqrt{3/\|\mathbf{P}\|_2}\}. \quad (3.8)$$

To prove these, observe that $\mathbf{X}_k \in \mathcal{C}_{\mathbf{P}}$ (and $\mathbf{Y}_k \in \mathcal{C}_{\mathbf{P}}$ for (3.4)) implies $\mathbf{X}_{k+1} \in \mathcal{C}_{\mathbf{P}}$ (and $\mathbf{Y}_{k+1} \in \mathcal{C}_{\mathbf{P}}$), so $\mathcal{C}_{\mathbf{P}}$ is an invariant submanifold under these iterations. Therefore, there exists an EVD

$P = Q\Lambda Q^T$ that also diagonalizes the iterates $X_k = Q\Lambda_k Q^T$ (and $Y_k = Q\Sigma_k Q^T$) and the matrix iteration can be decoupled into n scalar ones of the same form:

$$(3.2), (3.3) : \lambda_{k+1} = (1/2)(\lambda_k + \lambda/\lambda_k) : \lambda_k \rightarrow \sqrt{\lambda} \quad \text{for } \lambda_0 > 0 \quad (3.9)$$

$$(3.4) : \lambda_{k+1} = (1/2)(\lambda_k + 1/\sigma_k), \quad \sigma_{k+1} = (1/2)(\sigma_k + 1/\lambda_k) : \quad (3.10)$$

$$\lambda_k \rightarrow \sqrt{\lambda_0/\sigma_0}, \quad \sigma_k \rightarrow \sqrt{\sigma_0/\lambda_0} \quad \text{for } \lambda_0, \sigma_0 > 0$$

$$(3.5) : \lambda_{k+1} = (1/2)\lambda_k(3 - \lambda_k^2/\lambda) : \lambda_k \rightarrow 1/\sqrt{\lambda} \quad \text{for } 0 < \lambda_0 < \sqrt{3/\lambda} \quad (3.11)$$

Here $\lambda > 0$ is any diagonal entry of Λ , which is also an eigenvalue of P . Note that (3.10) can be reduced to two decoupled iterations of the same form as (3.9) due to the observation that $\lambda_k/\sigma_k \equiv \lambda_0/\sigma_0$. Since (3.9) is nothing but the well-known square root Newton iteration that solves $g(x) = x^2 - \lambda = 0$ while (3.11) is the Newton iteration for solving $g(x) = 1/x^2 - \lambda = 0$ (Philippe, 1987), we know both scalar iterations converge quadratically. Furthermore, this implies that the matrix iterations eqs. (3.2) to (3.5) also converge quadratically provided their initial values satisfy eqs. (3.6) to (3.8).

The results eqs. (3.6) to (3.8) only guarantee the iterations to converge analytically because \mathcal{C}_P as a submanifold of $\mathcal{M}_{n \times n}$ has no interior point and thus any structurally unconstrained perturbation to X_k such as numerical error, however small it is, will drive X_k away from \mathcal{C}_P . To ensure the stability of an iteration $X_{k+1} = F(X_k)$, we require the Fréchet derivative DF of the operator F at the limit X_* to have spectral radius $\rho(DF) \leq 1$. The Fréchet derivatives for eqs. (3.2) to (3.5) are summarized as

$$(3.2) : DF X = 0; \quad (3.12)$$

$$(3.3), (3.5) : DF X = (1/2)(X - \Lambda^{-1/2} X \Lambda^{1/2}); \quad (3.13)$$

$$(3.4) : DF(X, Y) = (1/2)(X - \Lambda^{1/2} Y \Lambda^{1/2}, Y - \Lambda^{-1/2} X \Lambda^{-1/2}). \quad (3.14)$$

In the above, (3.12) is a standard result for an exact Newton iteration, which must have $F(X) = X - (DG_X)^{-1}G(X)$. Hence, $(DF_{X_*})X = X - (DG_{X_*})^{-1}(DG_{X_*})X = X - X \equiv 0$. (3.13) and (3.14) are obtained by diagonalizing $P = Q\Lambda Q^T$ with a linear variable substitution $X' = Q^T X Q$. For example, for (3.13), originally we have $Y = DF X = (1/2)(X - P^{-1/2} X P^{1/2})$. After the variable substitution, $Y = DF X$ becomes $Y' = Q^T(DF(QX'Q^T))Q = (1/2)(X' - \Lambda^{-1/2} X' \Lambda^{1/2})$,

which yields (3.13). This simplification can be justified by the fact that the spectrum of $D\mathbf{F}$ is invariant under a linear variable substitution. Note that \mathbf{X}' is not diagonal in general because in stability analysis, we can no longer assume $\mathbf{X} \in \mathcal{C}_{\mathbf{P}}$.

Since each term in (3.13) and (3.14) has only one non-diagonal matrix factor, we can readily identify all the eigenvectors \mathbf{Z}_{ij} 's and eigenvalues μ_{ij} 's by observation:

$$(3.13) : \mathbf{Z}_{ij} = \mathbf{E}_{ij}, \quad \mu_{ij} = (1/2)(1 - \sqrt{\lambda_j/\lambda_i}); \quad (3.15)$$

$$(3.14) : \mathbf{Z}_{ij} = (\mathbf{E}_{ij}, 1/\sqrt{\lambda_i\lambda_j}\mathbf{E}_{ij}), \quad \mu_{ij} = 0, \quad (3.16)$$

$$\mathbf{Z}'_{ij} = (\mathbf{E}_{ij}, -1/\sqrt{\lambda_i\lambda_j}\mathbf{E}_{ij}), \quad \mu'_{ij} = 1.$$

Here $\mathbf{E}_{ij} \in \mathcal{M}_{n \times n}$ is the matrix having its (i, j) -th entry equal to 1 as the only nonzero entry. Therefore, if we want $\rho(D\mathbf{F}) = (1/2)(\sqrt{\kappa(\mathbf{P})} - 1) < 1$ for (3.15), we need the 2-norm condition number $\kappa(\mathbf{P})$ to satisfy $\kappa(\mathbf{P}) < 9$. This is very unsatisfactory because it puts a stringent constraint on the \mathbf{P} 's to which (3.3) and (3.5) are applicable. Philippe (1987) shows that for (3.5), a manual symmetrification step in the end of each iteration can alleviate this stability constraint to $\kappa(\mathbf{P}) < 17 + 6\sqrt{8}$, which is only a nonessential and minor improvement. This makes sense because $\mathcal{C}_{\mathbf{P}}$ requires being commutative with \mathbf{P} besides being SPD, so symmetrification itself does not retract a perturbed iterate back to $\mathcal{C}_{\mathbf{P}}$ in general, though may bring it closer.

On the other hand, (3.16) always has $\rho(D\mathbf{F}) = 1$. Note that (3.4) has the special feature that \mathbf{P} does not show up in the iteration at all and the converging point $(\mathbf{X}_*, \mathbf{Y}_*)$ can vary with the initial value $(\mathbf{X}_0, \mathbf{Y}_0)$ continuously. Indeed, the fixed points of (3.4) are not isolated but form an n^2 -dimensional submanifold corresponding to the n^2 degrees of freedom in $\mathbf{X}_0\mathbf{Y}_0^{-1}$ for specifying the initial values. Moreover, the n^2 eigenvectors associated with $\mu'_{ij} = 1$ in (3.16) span the tangent space of this fixed-point submanifold at \mathbf{X}_* . Any perturbation along these directions will neither grow nor diminish, but will simply remain constant and eventually reflect itself on a shift in the converging point, while perturbations along the eigenvectors for $\mu_{ij} = 0$ will decay quadratically.

Before moving on, we make several remarks on the practical use of eqs. (3.2) to (3.4). First, solving the Sylvester equation for $\Delta\mathbf{X}_k$ in (3.2) is nontrivial and expensive (Laub, 2005, sec. 13.3), which renders (3.2) of little practical value. Second, the extra constraint $\|\mathbf{X}_0\|_2 < \sqrt{3/\|\mathbf{P}\|_2}$ for the convergence of (3.5) indeed poses no difficulty because such an \mathbf{X}_0 is easy to find. Note

that $\|\mathbf{X}_0\|_2 < \sqrt{3/\|\mathbf{P}\|_2}$ can be replaced by a more stringent inequality $\|\mathbf{X}_0\| < \sqrt{3/\|\mathbf{P}\|}$ since $\|\mathbf{X}_0\|_2 = \rho(\mathbf{X}_0) \leq \|\mathbf{X}_0\|$ for any operator norm $\|\cdot\|$ and symmetric \mathbf{X}_0 . Hence, we can replace the 2-norm by another operator norm (e.g. 1-norm or ∞ -norm) that can be more easily calculated. A simple choice of \mathbf{X}_0 can be, for example, $\mathbf{X}_0 = (1/\sqrt{\|\mathbf{P}\|_1})\mathbf{I}$.

ii) Stabilizing Two Unstable SPD Square Root Iterations and Orthogonal Polar Factor

Next, we review two iterations for computing the orthogonal polar factor of an invertible $\mathbf{R} \in \mathcal{M}_{n \times n}$ (see Higham (1986a) and (Golub and Van Loan, 2013, sec. 9.4.3)):

$$\mathbf{X}_{k+1} = (1/2)(\mathbf{X}_k + \mathbf{X}_k^{-\text{T}}); \quad (3.17)$$

$$\mathbf{X}_{k+1} = \mathbf{X}_k + (1/2)\mathbf{X}_k(\mathbf{I} - \mathbf{X}_k^{\text{T}}\mathbf{X}_k), \quad (3.18)$$

where the initial value is set to $\mathbf{X}_0 = \mathbf{R}$ for both iterations. With the polar decomposition $\mathbf{X}_0 = \mathbf{Q}_0\mathbf{S}_0$, eqs. (3.17) and (3.18) is analytically equivalent to using the same iterations with initial value $\mathbf{X}_0 = \mathbf{S}_0$ to compute the square root of \mathbf{I} . Note that eqs. (3.17) and (3.18) are almost the same as eqs. (3.3) and (3.5) with $\mathbf{P} = \mathbf{I}$ except for the presence of an extra transpose. Analytically this makes no difference when the initial value \mathbf{X}_0 is SPD since all the iterates will be symmetric. However, surprisingly it turns out that this extra transpose makes a huge difference in numerical stability because it indeed lifts the stringent $\kappa(\mathbf{P}) < 9$ stability constraint for eqs. (3.3) and (3.5).

We will expound this significant improvement by generalizing eqs. (3.17) and (3.18) to the following multi-purpose (SPD square root or orthogonal polar factor) iterations:

$$\mathbf{X}_{k+1} = (1/2)(\mathbf{X}_k + \mathbf{X}_k^{-\text{T}}\mathbf{P}); \quad (3.19)$$

$$\mathbf{X}_{k+1} = \mathbf{X}_k + (1/2)\mathbf{X}_k(\mathbf{I} - \mathbf{X}_k^{\text{T}}\mathbf{P}\mathbf{X}_k), \quad (3.20)$$

where \mathbf{P} is some SPD matrix. eqs. (3.19) and (3.20) can be viewed as the counterparts of eqs. (3.3) and (3.5) for computing the square root of \mathbf{P} in the sense of $\mathbf{X}^{\text{T}}\mathbf{X} = \mathbf{P}$ rather than $\mathbf{X}^2 = \mathbf{P}$. Moreover, (3.20) also has a nice interpretation as a successive first order approximation to $\mathbf{X}^{\text{T}}\mathbf{X} - \mathbf{P} = \mathbf{0}$, which is shown in appendix B.

Suppose the initial value \mathbf{X}_0 is invertible and has polar decomposition $\mathbf{X}_0 = \mathbf{Q}_0\mathbf{S}_0$ (or $\mathbf{X}_0^{\text{T}} =$

$\mathbf{Q}_0^T \mathbf{S}_0$). Then we have the following convergence properties:

$$(3.19) : \mathbf{X}_k \rightarrow \mathbf{Q}_0 \mathbf{P}^{1/2} \quad \text{for } \mathbf{X}_0 = \mathbf{Q}_0 \mathbf{S}_0, \mathbf{S}_0 \in \mathcal{C}_P; \quad (3.21)$$

$$(3.20) : \mathbf{X}_k \rightarrow \mathbf{P}^{-1/2} \mathbf{Q}_0 \quad \text{for } \mathbf{X}_0 = \mathbf{S}_0 \mathbf{Q}_0, \mathbf{S}_0 \in \mathcal{C}_P, \|\mathbf{X}_0\|_2 < \sqrt{3/\|\mathbf{P}\|_2}. \quad (3.22)$$

In other words, \mathbf{X}_k converges to a particular (inverse) square root of \mathbf{P} that shares the same orthogonal polar factor with \mathbf{X}_0 . In particular, if we set $\mathbf{P} = \mathbf{I}$, \mathbf{X}_k will converge to the orthogonal polar factor \mathbf{Q}_0 of \mathbf{X}_0 , while if we instead set $\mathbf{Q}_0 = \mathbf{I}$ (i.e. $\mathbf{X}_0 \in \mathcal{C}_P$), \mathbf{X}_k will converge to the (inverse) SPD square root of \mathbf{P} . Hence, these two multi-purpose iterations reflect the inherent connections between matrix square root and polar decomposition and unify the two on a algorithmic level.

Using the same methodology for stability analysis as before, we obtain the following Fréchet derivative at \mathbf{X}_* for both eqs. (3.19) and (3.20):

$$\mathbf{D}\mathbf{F}\mathbf{X} = (1/2)(\mathbf{X} - \mathbf{\Lambda}^{-1/2} \mathbf{X}^T \mathbf{\Lambda}^{1/2}), \quad (3.23)$$

where the diagonal $\mathbf{\Lambda}$ contains the eigenvalues of \mathbf{P} . Again, by observation, we obtain all the eigenvalues and eigenvectors for the linear operator $\mathbf{D}\mathbf{F}$ as:

$$\begin{aligned} \mathbf{Z}_{ij} &= \mathbf{E}_{ij} + \sqrt{\lambda_i/\lambda_j} \mathbf{E}_{ji}, \quad \mu_{ij} = 0, \quad i \leq j, \\ \mathbf{Z}_{ij} &= \mathbf{E}_{ij} - \sqrt{\lambda_i/\lambda_j} \mathbf{E}_{ji}, \quad \mu_{ij} = 1, \quad i > j. \end{aligned} \quad (3.24)$$

Therefore, instead of the stringent constraint $\kappa(\mathbf{P}) < 9$ for eqs. (3.3) and (3.5) due to (3.15), now we have $\rho(\mathbf{D}\mathbf{F}) \equiv 1$ regardless of $\kappa\mathbf{P}$. Moreover, the fixed points are no longer isolated and they form a $n(n-1)/2$ -dimensional submanifold corresponding to the unitary freedom encoded by \mathbf{Q}_0 . The $n(n-1)/2$ eigenvectors associated with eigenvalue 1 span the tangent space of this fixed-point submanifold at \mathbf{X}_* and any perturbation orthogonal to this tangent space decays quadratically to 0.

Again the extra initial value constraint $\|\mathbf{X}_0\|_2 < \sqrt{3/\|\mathbf{P}\|_2}$ for (3.20) poses no practical difficulty. The trade-off between (3.19) and (3.20) lies in the relative cost between matrix inversion $\mathbf{A}^{-1}\mathbf{B}$ and multiplication $\mathbf{A}\mathbf{B}$ for square \mathbf{A} and \mathbf{B} . (3.19) costs one inversion while (3.20) costs three

multiplications. In practice, computing $\mathbf{A}^{-1}\mathbf{B}$ is about 1.5 times the cost of computing \mathbf{AB} (based on testing in MATLAB and Higham and Schreiber (1990)), so overall eqs. (3.19) and (3.20) are equally good alternatives to the SVD/EVD approach for computing the SPD square root or the orthogonal polar factor.

Finally, we emphasize that unlike eqs. (3.2) to (3.5) which aim to solve $\mathbf{X}^2 = \mathbf{P}$, eqs. (3.19) and (3.20) target at solving $\mathbf{X}^T\mathbf{X} = \mathbf{P}$. Consequently, eqs. (3.19) and (3.20) will not yield the major square root of a general \mathbf{P} with eigenvalues not on the negative real axis.

The various new algorithms recommended for solving the minimum-correction second-moment matching problem in theorem 2.1 are summarized in table 3.1.

Algorithm	Description
\mathbf{A}_* -EVD	Use (3.1) for \mathbf{A}_* with $\sqrt{\cdot}$ computed by Cholesky factorization and $(\cdot)^{1/2}$ computed by EVD (see section 3.2.1).
\mathbf{A}_* -NtSqr	Use (3.1) for \mathbf{A}_* with $\sqrt{\cdot}$ computed by Cholesky factorization and $(\cdot)^{1/2}$ computed by the stablized Newton iteration (3.19) or (3.20).
$\tilde{\mathbf{U}}_*$ -NtPD	Perform Householder QR factorization $\mathbf{U} = \mathbf{V}\mathbf{R}$. Compute the orthogonal polar factor \mathbf{Q} of $\mathbf{R}\sqrt{\tilde{\mathbf{P}}}^T$ using the stablized Newton iteration (3.19) or (3.20). Form $\tilde{\mathbf{U}}_* = \mathbf{V}\mathbf{Q}\sqrt{\tilde{\mathbf{P}}}$.

Table 3.1: Recommended algorithms for computing the minimum-correction second-moment matching solution $\tilde{\mathbf{U}}_*$ in theorem 2.1. $\tilde{\mathbf{U}}_*$ -based algorithms avoid forming $\mathbf{P} = \mathbf{U}^T\mathbf{U}$ and directly operate on \mathbf{U} .

3.3 The Special Case of Orthonormalization

Now we turn to the special case of orthonormalization, i.e. $\tilde{\mathbf{P}} = \mathbf{I}$, where (3.1) reduces to $\mathbf{A}_* = \mathbf{P}^{-1/2}$ and $\mathbf{U} = \tilde{\mathbf{U}}_*\mathbf{A}_*^{-1}$ becomes the polar decomposition. In this case, Higham (1988) provides an a priori bound for $\|\tilde{\mathbf{U}}_* - \mathbf{U}\|_F$:

$$\|\mathbf{U}^T\mathbf{U} - \mathbf{I}\|_F/(\|\mathbf{U}\|_2 + 1) \leq \|\tilde{\mathbf{U}}_* - \mathbf{U}\|_F \leq \|\mathbf{U}^T\mathbf{U} - \mathbf{I}\|_F. \quad (3.25)$$

When $\|\mathbf{U}\|_2$ is of order 1 or smaller, the above yields a tight estimate of the order of magnitude of $\|\tilde{\mathbf{U}}_* - \mathbf{U}\|_F$. One common scenario of such is with \mathbf{U} being nearly orthogonal, i.e. $\mathbf{P} = \mathbf{U}^T \mathbf{U} \approx \mathbf{I}$, when, for example, a path on a Stiefel manifold is tracked numerically for subspace tracking or optimization with orthogonality constraints.

For orthonormalization, there are many widely-used algorithms that yield an orthogonal matrix not being the closest to the original \mathbf{U} . We clarify next how they are related to the minimum-correction solution.

Comparing with QR-Based Orthonormalization If we only seek an orthogonal $\tilde{\mathbf{U}}$ close to \mathbf{U} , but not necessarily the closest one, we can stop after the QR factorization step $\mathbf{U} = \mathbf{V}\mathbf{R}$ in section 3.1 and take $\tilde{\mathbf{U}} = \mathbf{V}$. This corresponds to the Cholesky decomposition $\mathbf{P} = \mathbf{R}^T \mathbf{R}$. The continuity of the Cholesky factor \mathbf{R} with respect to \mathbf{P} (Schatzman, 2002, sec. 12.1.3) guarantees that $\mathbf{P} \approx \mathbf{I}$ implies $\mathbf{R} \approx \mathbf{I}$ and thus $\mathbf{V} \approx \mathbf{U}$. More precisely,

$$\|\tilde{\mathbf{U}}_* - \mathbf{U}\|_F \leq \|\mathbf{V} - \mathbf{U}\|_F \leq \frac{1 + \|\mathbf{U}\|_2}{\sqrt{2}(1 - \|\mathbf{U}^T \mathbf{U} - \mathbf{I}\|_2)} \|\tilde{\mathbf{U}}_* - \mathbf{U}\|_F, \quad (3.26)$$

which is credited to Jiguang Sun by Higham (1997), bounds the closeness of the orthogonal QR factor \mathbf{V} to \mathbf{U} by a multiple of the minimum $\|\tilde{\mathbf{U}}_* - \mathbf{U}\|_F$ (see also (Chandrasekaran and Ipsen, 1994, sec. 4)). In particular, we have approximately $\|\mathbf{V} - \mathbf{U}\|_F \leq \sqrt{2}\|\tilde{\mathbf{U}}_* - \mathbf{U}\|_F$ when $\mathbf{P} \approx \mathbf{I}$.

This closeness property of the QR factorization stems from the nested subspace correspondence between \mathbf{V} and \mathbf{U} , meaning that the first $k = 1, \dots, n$ columns of \mathbf{V} and \mathbf{U} span the same subspace in \mathbb{R}^m , which constrains \mathbf{V} to be close to \mathbf{U} . This also implies that when \mathbf{U} 's columns are permuted, the corresponding orthogonal QR factor \mathbf{V} cannot be obtained by applying the same permutation to the original one, unlike the orthogonal polar factor $\tilde{\mathbf{U}}_*$. Such column-order-dependence could be advantageous because some columns of \mathbf{U} might be numerically more accurate than the others in certain cases. Hence, swapping the more accurate columns to the front allows us to shift the modification of \mathbf{U} towards its less accurate part.

As mentioned before, the algorithm of choice for QR factorization is commonly Householder reflection (Trefethen and Bau III, 1997, ch. 10). The only caveat is that the Householder QR factorization

in general does not ensure the positivity of the upper triangular \mathbf{R} 's diagonal entries. However, this can be easily fixed a posteriori by flipping the sign of the columns of \mathbf{V} corresponding to the negative diagonal entries of \mathbf{R} . On the other hand, in the case of nearly-orthogonal \mathbf{U} , \mathbf{U} is very well-conditioned so the stability of Householder QR factorization over modified Gram-Schmidt or even classical Gram-Schmidt becomes uncritical, as will be seen in chapter 5.

Comparing with Plain SVD Orthonormalization (Ueckermann et al., 2011, app. B) proposes a plain SVD algorithm where the left singular vectors \mathbf{V} of an SVD $\mathbf{U} = \mathbf{V}\mathbf{\Sigma}\mathbf{Z}^T$ is directly taken as $\tilde{\mathbf{U}}$. The prime issue with this approach is the non-uniqueness of an SVD. The singular vectors associated with the same (or numerically close) singular value are only determined up to an orthogonal transform, which can flip signs, permute the order and even rotate the vectors. This is especially problematic when \mathbf{U} is nearly orthogonal because all its singular values will cluster around 1. In this case, $\|\mathbf{V} - \mathbf{U}\|_F$ can be of order 1 even if \mathbf{U} is nearly orthogonal. To remedy this issue, Feppon and Lermusiaux (2018b) suggest a gradient descent algorithm for orthonormalization. Here we analyze it in appendix C and conclude that due to its linear convergence and stringent stability constraint, it is unattractive compared to the algorithms in section 3.2.

Indeed, as indicated in section 3.2.1, the plain SVD approach can be easily modified to directly yield the polar decomposition $\mathbf{U} = (\mathbf{V}\mathbf{Z}^T)(\mathbf{Z}\mathbf{\Sigma}\mathbf{Z}^T)$, where $\tilde{\mathbf{U}}_* = \mathbf{V}\mathbf{Z}^T$ is the minimum-correction solution. Note how this extra \mathbf{Z}^T factor eliminates the unitary freedom causing all the undesired sign flips, permutation, and rotation.

Chapter 4

Benchmarks of the Algorithms' Performance

In this section, we benchmark the performance of the various new minimum-correction second-moment matching algorithms introduced in chapter 3 for the special case of orthonormalization (i.e. $\tilde{\mathbf{P}} = \mathbf{I}$) and compare them to existing algorithms. All the algorithms to be tested are described in table 4.1. Note that PD-NtSqr, PD-NtISqr and PD-NtPD are based on our new multi-purpose Newton iteration (3.19) and (3.20), while others are existing algorithms from literature. Comparing table 4.1 to table 3.1, here PD-EVD corresponds to \mathbf{A}_* -EVD in table 3.1, PD-NtSqr and PD-NtISqr to \mathbf{A}_* -NtSqr, while PD-NtPD corresponds to $\tilde{\mathbf{U}}_*$ -NtPD.

For all algorithms, the fixed-point iteration terminates once $\|\mathbf{X}_k - \mathbf{X}_{k-1}\| < 10^{-14}\|\mathbf{X}_{k-1}\|$ is satisfied or the maximum number of iterations is reached. The numerical tests are performed under the MATLAB[®] environment on a machine with an Intel[®] Core[™] i7-4702MQ processor and 16GB of memory.

The algorithms are tested in three cases. The first is with a $10^6 \times 10^2$ matrix \mathbf{U} (m -by- n , $m \gg n$) of a small condition number $\kappa = 1.5$, so the tall thin \mathbf{U} is nearly orthogonal. We form $\mathbf{U} = \mathbf{U}_0 \mathbf{\Lambda} \mathbf{Q}_0$ by random generation of an orthogonal m -by- n \mathbf{U}_0 and orthogonal n -by- n \mathbf{Q}_0 . A random orthogonal matrix is obtained by applying QR Householder to a random matrix whose entries are independent

Algorithm	Description
SVD-EVD	Plain SVD ortho. (see section 3.3) by EVD of \mathbf{P} .
PD-EVD	PD ortho. by EVD of \mathbf{P} (see section 3.2.1).
PD-SVD*	PD ortho. by SVD of \mathbf{U} (see section 3.2.1).
PD-NtSqr	PD ortho. by computing $\mathbf{P}^{1/2}$ using Newton iteration (3.19).
PD-NtISqr	PD ortho. by computing $\mathbf{P}^{-1/2}$ using Newton iteration (3.20).
PD-GD	PD ortho. by computing $\mathbf{P}^{-1/2}$ using gradient descent (C.2).
PD-NtPD*	PD ortho. by Householder QR $\mathbf{U} = \mathbf{V}\mathbf{R}$ and computing the orthogonal polar factor of \mathbf{R} using Newton iteration (3.19).
QR-CGS*	QR ortho. by classical Gram-Schmidt.
QR-MGS*	QR ortho. by modified Gram-Schmidt.
QR-Chol	QR ortho. by Cholesky decomposition of \mathbf{P} .
QR-HH*	QR ortho. by Householder QR $\mathbf{U} = \mathbf{V}\mathbf{R}$.

Table 4.1: Description of the various orthonormalization algorithms benchmarked. Here \mathbf{U} is to be orthonormalized and $\mathbf{P} = \mathbf{U}^T \mathbf{U}$ is its Gram matrix. The algorithms that do not form \mathbf{P} and directly operate on \mathbf{U} are marked by a “*” in their names.

and uniformly distributed over $[-1, 1]$. \mathbf{A} is a diagonal matrix of condition number κ and of diagonal entries decaying geometrically to 1. Each algorithm is tested using the same 10 samples of \mathbf{U} and the average performance metrics are listed in table 4.2. For the second and third cases, the same tests are performed, but for a square well-conditioned \mathbf{U} and a rectangular ill-conditioned \mathbf{U} , respectively. Results are in table 4.3.

Overall, tables 4.2 and 4.3 confirm several key messages:

i The plain SVD algorithm does not promote closeness between $\tilde{\mathbf{U}}$ and \mathbf{U} , as is verified by $\|\tilde{\mathbf{U}} - \mathbf{U}\|$ of SVD-EVD being much larger than that of the QR-based and of the new minimum-correction (PD-based) algorithms, see table 4.2.

ii The new PD-based algorithms achieve the minimum correction as is verified by the agreement between their $\|\tilde{\mathbf{U}} - \mathbf{U}\|$ values and the analytical one computed using (2.11) as $\|\tilde{\mathbf{U}}_\star - \mathbf{U}\| = (\text{tr} \mathbf{I} + \text{tr} \mathbf{P} - 2\text{tr} \mathbf{P}^{1/2})^{1/2} = (n + \|\mathbf{U}\|_{\text{F}}^2 - 2\|\mathbf{U}\|_\star)^{1/2}$, where $\|\cdot\|_\star$ is the nuclear norm, being the sum of singular values. The symmetry of \mathbf{A}_\star is verified by the machine epsilon values of $\|\mathbf{A} - \mathbf{A}^T\|$.

$m = 10^6, n = 10^2, \kappa(\mathbf{U}) = 1.5$				
	$\ \tilde{\mathbf{U}}^T \tilde{\mathbf{U}} - \mathbf{I}\ $	$\ \tilde{\mathbf{U}} - \mathbf{U}\ $	$\ \mathbf{A} - \mathbf{A}^T\ $	(#Ite) Time
SVD-EVD	2e-14	16	12	0.6s
PD-EVD	3e-14	2.75	1e-15	0.6s
PD-SVD*	3e-14	2.75	1e-15	5s
PD-NtSqr	8e-15	2.75	8e-16	(6) 2.3s
PD-NtISqr	9e-15	2.75	1e-15	(8) 0.6s
PD-GD	8e-13	2.75	3e-15	(131) 0.6s
PD-NtPD*	1e-14	2.75	1e-15	(7) 5s
QR-CGS*	6e-15	3.04	1.9	38s
QR-MGS*	6e-15	3.04	1.9	38s
QR-Chol	7e-15	3.04	1.9	2s
QR-HH*	1e-14	3.04	1.9	4s

Table 4.2: Performance of the various algorithms for orthonormalization in the case of well-conditioned (nearly orthogonal) rectangular matrices (m -by- n , $m \gg n$). Here \mathbf{A} is the linear transform such that $\tilde{\mathbf{U}} = \mathbf{U}\mathbf{A}$. The Frobenius norm is exclusively used. For the algorithms based on a fixed-point iteration, the average number of iterations is recorded in the parenthesis in front of the run time. In this case, we have $\|\mathbf{U}^T \mathbf{U} - \mathbf{I}\| = 6.5$ for all 10 repetitions, which can be compared to $\|\tilde{\mathbf{U}}^T \tilde{\mathbf{U}} - \mathbf{I}\|$.

	$m = n = 2 \times 10^3, \kappa(\mathbf{U}) = 1.5$			$m = 10^6, n = 10^2, \kappa = 10^6$	
	$\ \tilde{\mathbf{P}} - \mathbf{I}\ $	$\ \tilde{\mathbf{U}} - \mathbf{U}\ $	(#Ite) Time	$\ \tilde{\mathbf{P}} - \mathbf{I}\ $	(#Ite) Time
PD-EVD	3e-13	12.3	2s	4e-5	0.6s
PD-SVD*	4e-13	12.3	4s	3e-14	5s
PD-NtSqr	4e-14	12.3	(7) 3s	4e-5	(100) 2s
PD-NtInv	6e-14	12.3	(9) 6s	4e-5	(100) 0.6s
PD-GD	3e-11	12.3	(422) 280s	9	(1e4) 3s
PD-NtPD*	7e-14	12.3	(9) 4s	2e-14	(26) 5s
QR-CGS*	8e-14	13.6	13s	6e-5	38s
QR-MGS*	6e-14	13.6	13s	2e-10	38s
QR-Chol	3e-14	13.6	0.3s	3e-5	2s
QR-HH*	7e-14	13.6	0.4s	1e-14	4s

Table 4.3: Performance of the algorithms on well-conditioned square matrices ($\kappa(\mathbf{U}) = 1.5$, $\|\mathbf{U}^T \mathbf{U} - \mathbf{I}\| = 29$) and on ill-conditioned rectangular matrices ($\kappa(\mathbf{U}) = 10^6$, $\|\mathbf{U}^T \mathbf{U} - \mathbf{I}\| = 1.5 \times 10^{12}$), respectively. Notations are as those in table 4.2.

The $\|\tilde{U} - U\|$ values of the QR-based algorithms are larger than $\|\tilde{U}_* - U\|$, but within the bound (3.26).

iii For ill-conditioned matrices, the algorithms based on forming $P = U^T U$ suffer from poor accuracy due to squaring the condition number. This is revealed by the large $\|\tilde{P} - I\|$ values of the algorithms without a “*” in their names, see table 4.3. Among them, those based on a fixed-point iteration have their solutions stagnate with large orthogonality errors after only a few iterations. The iterates neither converge nor blow up and just bounce around, even after as many as 10^5 iterations. CGS and MGS are known to suffer from the same orthogonality issue in this case. Hence, the only algorithms that can still reliably orthonormalize U are those based on performing Householder QR or SVD on U .

iv For tall thin matrices, the algorithms based on forming $P = U^T U$ tend to be more efficient when U is well-conditioned, because other than computing the small P in the beginning and applying A to U in the end, they only deal with small n -by- n matrices. This strategy avoids performing QR factorization or SVD to a large m -by- n matrix, which could be relatively expensive in this case.

Chapter 5

Application to Matrix Differential Equations that Preserve Orthogonality

A matrix differential equation

$$\mathrm{d}_t U = F(U) \tag{5.1}$$

for $U(t) \in \mathcal{M}_{m \times n}$ is said to preserve orthogonality if $U^T U = I$ at $t = 0$ implies $U^T U = I$ at all $t > 0$. We encounter such equations in applications such as subspace tracking (Brockett, 1991; Lermusiaux, 1997, 2001; Koch and Lubich, 2007; Bergemann et al., 2009; Feppon and Lermusiaux, 2018a), where we track the time evolution of an n -dimensional subspace of \mathbb{R}^m , which may come from a discretization of an infinite-dimensional dynamical system, by evolving an orthonormal basis with its base vectors forming U 's columns. In this section, we showcase our new minimum-correction second-moment matching algorithms in the special case of $\tilde{P} = I$ for re-orthonormalization required in solving such equations. Before the examples, we analyze the error due to the loss of numerical orthogonality and clarify the role played by re-orthonormalization.

Although eq. (5.1) preserves orthogonality, its time-discretized counterpart $\hat{U}_{n+1} = \hat{F}(\hat{U}_n)$ in general does not (Dieci et al., 1994; Higham, 1996, 1997; Ueckermann et al., 2013). Here, \hat{U}_n is a

numerical approximation of $U_n = U(t_n)$ and the map \hat{F} stems from a particular time-marching scheme, which is said to preserve orthogonality if $\hat{U}_n^T \hat{U}_n = I \implies \hat{U}_{n+1}^T \hat{U}_{n+1} = I$. According to Dieci et al. (1994) and Higham (1996), the only known family of orthogonality-preserving time-marching schemes is the one of Gauss-Legendre Runge-Kutta schemes. Unfortunately, such schemes are difficult to implement, expensive to use and not widely available in numerical software (Higham, 1997). Moreover, they only preserve orthogonality for certain F , which renders them even less attractive. Therefore, a more practical strategy is to integrate (5.1) by an arbitrary scheme and re-orthonormalize the \hat{U}_n 's at the end of each time step.

We denote by \hat{U} the numerical solution at a particular time obtained by some time-marching scheme without re-orthonormalization, and by $\Delta U = \hat{U} - U$ the time-marching error. If we choose $\|\cdot\|$ as the matrix 2-norm, we have

$$\|\hat{U}^T \hat{U} - I\| = \|(U + \Delta U)^T (U + \Delta U) - I\| \leq 2 \|\Delta U\| \|U\| + \|\Delta U\|^2,$$

since $U^T U = I$. If the scheme is p -th order accurate, i.e. $\|\Delta U\| = O((\Delta t)^p)$, since $\|U\| = O(1)$, we have $\|\hat{U}^T \hat{U} - I\| = O((\Delta t)^p)$, which is of the same order as the time-marching error.¹ This seems to imply that as long as the time-marching scheme is accurate enough, there is no need to worry about the orthogonality error. However, the fact that (5.1) preserves orthogonality only implies that the Stiefel manifold $\mathcal{S}_I = \{U : U^T U = I\}$ is invariant under (5.1), but \mathcal{S}_I may not be a stable manifold. Therefore, any deviation of U from \mathcal{S}_I may not stay bounded under (5.1). It might even hit an unstable direction and diverge exponentially. In addition, in many cases, the derivation and thus the validity of (5.1) itself relies on U remaining orthogonal, so a violation of this condition may unexpectedly ruin the numerical solution.

The above analysis justifies re-orthonormalizing \hat{U} after each time step. This procedure should ideally not undermine the convergence order of the time-marching scheme or introduce significant numerical artifacts. Hence, we need to analyze how re-orthonormalization affects the numerical errors.

Suppose $\tilde{U} \in \mathcal{S}_I$ is the orthonormalized solution obtained from \hat{U} . To maintain the order of

¹ It is not really “one order higher than the temporal discretization errors” as claimed in (Ueckermann et al., 2013, sec. 5.2), where the one-order-higher effect of the *local* truncation error was misinterpreted.

convergence, it suffices to have $\|\tilde{\mathbf{U}} - \hat{\mathbf{U}}\| \leq C\|\hat{\mathbf{U}} - \mathbf{U}\|$ for some $C > 0$ because then by triangle inequality, we have $\|\tilde{\mathbf{U}} - \mathbf{U}\| \leq (C+1)\|\hat{\mathbf{U}} - \mathbf{U}\|$. This can be satisfied if we choose $\tilde{\mathbf{U}}$ to be $\tilde{\mathbf{U}}_*$, the orthogonal projection of $\hat{\mathbf{U}}$ onto $\mathcal{S}_{\mathbf{I}}$, because, as $\mathbf{U} \in \mathcal{S}_{\mathbf{I}}$, corollary 2.4 implies $\|\tilde{\mathbf{U}}_* - \hat{\mathbf{U}}\| \leq \|\hat{\mathbf{U}} - \mathbf{U}\|$ and thus $\|\tilde{\mathbf{U}}_* - \mathbf{U}\| \leq 2\|\hat{\mathbf{U}} - \mathbf{U}\|$. Indeed, since $\|\hat{\mathbf{U}} - \mathbf{U}\|$ is small, $\tilde{\mathbf{U}}_*$, $\hat{\mathbf{U}}$, and \mathbf{U} are almost the vertices of a right triangle (“almost” because $\mathcal{S}_{\mathbf{I}}$ is not flat), so in practice, $\|\tilde{\mathbf{U}}_* - \mathbf{U}\| \leq \|\hat{\mathbf{U}} - \mathbf{U}\|$. Numerical examples in Dieci et al. (1994) show that in some cases $\|\tilde{\mathbf{U}}_* - \mathbf{U}\|$ can even be smaller than $\|\hat{\mathbf{U}} - \mathbf{U}\|$ by an order of magnitude. Another viable option for $\tilde{\mathbf{U}}$ is the orthogonal QR factor \mathbf{V} of $\hat{\mathbf{U}}$ because (3.26) guarantees $\|\mathbf{V} - \hat{\mathbf{U}}\| \leq \sqrt{2}\|\tilde{\mathbf{U}}_* - \hat{\mathbf{U}}\| \leq \sqrt{2}\|\hat{\mathbf{U}} - \mathbf{U}\|$.

To showcase the new algorithms and demonstrate the above results, we consider the stochastic Lorenz-96 system (Lorenz, 1996) (the index i is circular, i.e. $x_i = x_{i+m}$),

$$\mathrm{d}_t x_i = (x_{i+1} - x_{i-2})x_{i-1} - x_i + F, \quad i = 1, \dots, m. \quad (5.2)$$

This toy model mimics atmospheric convection along a mid-latitude circle. We set $m = 40$ and $F = 8$ which lead to chaotic dynamics. In all simulations, the initial condition is $x_i = 0$ for $i \neq 1$ while x_1 is uniformly distributed over $[-0.01, 0.01]$. The classical explicit RK-4 scheme is used for time integration.

To solve (5.2), we employ the dynamically orthogonal (DO) equations (Sapsis and Lermusiaux, 2009). They approximate the random vector $\mathbf{u}(t; \omega) = [x_1, \dots, x_m]^T \in \mathbb{R}^m$ by a low rank expansion $\mathbf{u}(t; \omega) = \bar{\mathbf{u}}(t) + \sum_{i=1}^n \phi_i(t; \omega) \mathbf{u}_i(t) = \bar{\mathbf{u}} + \mathbf{U} \boldsymbol{\phi}$ where $\bar{\mathbf{u}}$ is the mean, $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_n] \in \mathcal{M}_{m \times n}$ contains the orthonormal base vectors and $\boldsymbol{\phi} = [\phi_1, \dots, \phi_n]^T$ contains the random scalar coefficients. The evolution of all these components is governed by a dynamical system, the DO equations, that preserves and whose validity is predicated on $\mathbf{U}^T \mathbf{U} \equiv \mathbf{I}$. Such techniques are useful for uncertainty quantification in high-dimensional systems governed by time-dependent stochastic PDEs (Feppon and Lermusiaux, 2018b; Subramani, 2018; Subramani and Lermusiaux, 2019). Related works include the dynamic low-rank approximation of a time-varying matrix (Koch and Lubich, 2007), geometric analyses of the DO equations based on matrix manifold theory (Feppon and Lermusiaux, 2018a, 2019) and dynamically bi-orthogonal equations (Cheng et al., 2013).

In fig. 5-1, we first show the time evolution of the orthogonality error $\|\mathbf{U}^T \mathbf{U} - \mathbf{I}\|_{\max}$ ($\|\mathbf{A}\|_{\max} \triangleq \max\{A_{ij}\}$) for the numerical integration of the DO equations without any re-orthonormalization.




Figure 5-1: Orthogonality errors in the absence of re-orthonormalization. The error saturates at a higher level with larger time step Δt . When $\Delta t = 0.05$, the orthogonality is completely lost for $n = 10$ (left), while the solution even blows up for $n = 20$ (right).

Results are shown for $n = 10, 20$ and three time step sizes. As we can see, the orthogonality error grows with Δt since a larger Δt implies a larger discrepancy between the time-discrete system and the originally orthogonality-preserving continuous one. Moreover, when $\Delta t = 0.005$ and 0.01 , $\|U^T U - I\|_{\max}$ tends to saturate at a low level ($< 10^{-3}$) that corresponds to only a negligible deviation from orthogonality. However, when Δt increases to 0.05 , $\|U^T U - I\|_{\max}$ can either level off at a magnitude of 1 for $n = 10$ that signals a complete loss of orthogonality or even drive the simulation unstable and cause the solution to blow up for $n = 20$. This exemplifies the necessity of re-orthonormalization.

Note that in the expansion $\mathbf{u} = \bar{\mathbf{u}} + \mathbf{U}\phi$, there is also the unitary freedom $\mathbf{U}\phi = (\mathbf{U}\mathbf{Q})(\mathbf{Q}^T\phi)$ for any orthogonal \mathbf{Q} , under which all realizations $\mathbf{u}(t; \omega)$ are unchanged. Therefore, if we apply such a \mathbf{Q} at some time t and use $\mathbf{U}\mathbf{Q}$ and $\mathbf{Q}^T\phi$ as the initial values of the modes and coefficients for the subsequent time integration, the future evolution of $\mathbf{u}(t; \omega)$ will remain the same. However, this breaks the time continuity of \mathbf{U} and ϕ at time t , as illustrated in fig. 5-2, where the SVD-EVD and PD-EVD (see table 4.1) algorithm are compared. Due to the non-uniqueness of an SVD, SVD-EVD not only removes the orthogonality error, but also applies to \mathbf{U} some \mathbf{Q} that may not at all be close to \mathbf{I} . As a result, we lose the time continuity of the modes and coefficients as indicated by the “random” pattern in the left half of fig. 5-2. In contrast, PD-EVD minimizes such numerical artifacts. Maintaining time continuity is especially important when we use multi-step

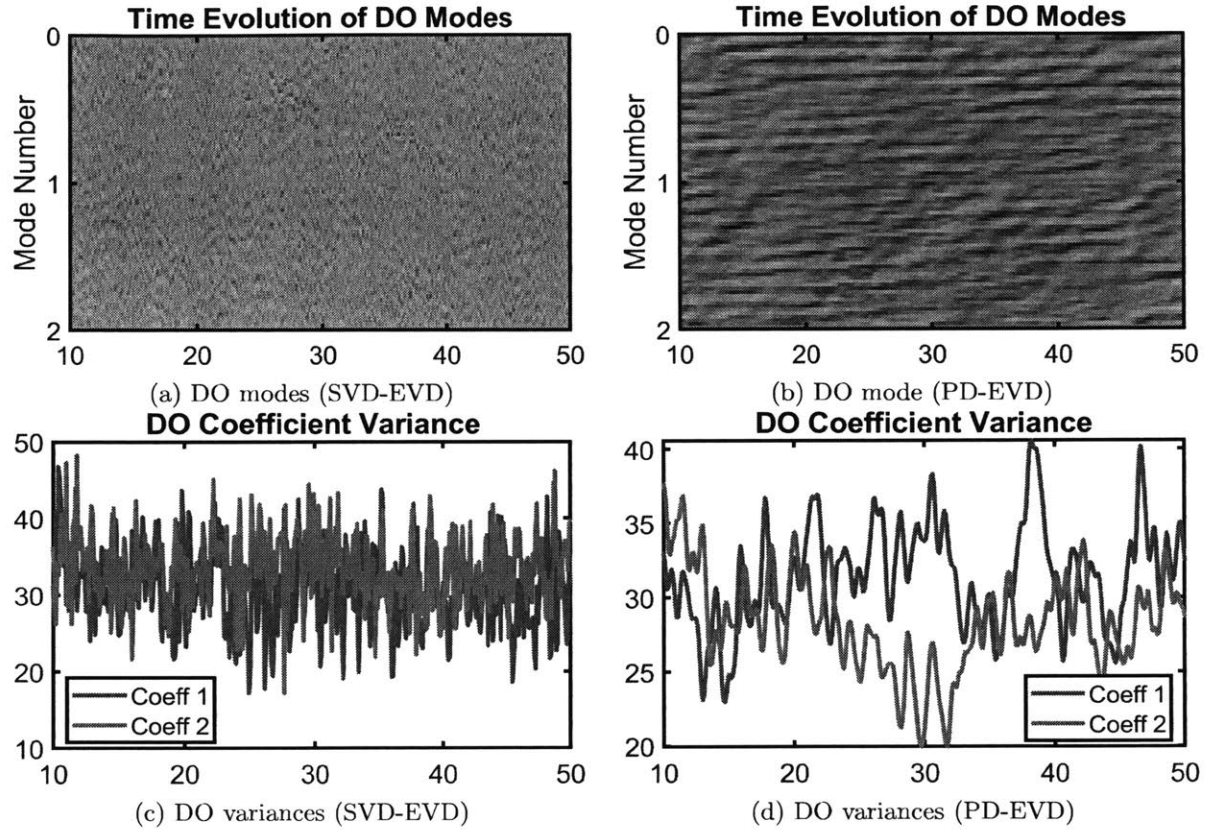


Figure 5-2: Comparison between orthonormalization using SVD-EVD (left) with that using polar decomposition PD-EVD (right). The evolution (from $t = 10$ to 50 , $\Delta t = 0.05$) of the first two DO modes $[\mathbf{u}_1^T, \mathbf{u}_2^T]^T$ stacked together are shown on the top and the evolution of their coefficient variances $\text{Var}(\phi_1)$ and $\text{Var}(\phi_2)$ on the bottom.

time-marching schemes² because their validity relies on the smooth evolution of \mathbf{U} and ϕ across different time steps. A “random” \mathbf{Q} applied at each time step destroys this smoothness and can cause the time integration to diverge.

Finally we show in fig. 5-3 how the orthogonality error can be controlled by PD-EVD, PD-NtSqr, PD-NtlSqr, QR-Chol and SVD-EVD for re-orthonormalization. These algorithms were described in table 4.1. They are the best choices in this case based on the discussion in chapter 4 because the $\hat{\mathbf{U}}$ to be re-orthonormalized after each time step is a tall thin matrix and is already nearly orthogonal. Here we exclusively use $\Delta t = 0.05$ and $n = 10$ for all simulations. The left of fig. 5-3 indicates that the orthogonality errors are kept at almost machine epsilon. The two EVD-based algorithms have slightly larger errors, most likely due to the convergence criterion for the built-in EVD subroutine, but this is not essential in practice. The right plot shows the ratio of the actual correction $\|\tilde{\mathbf{U}} - \hat{\mathbf{U}}\|$ to the theoretical minimum $\|\tilde{\mathbf{U}}_* - \hat{\mathbf{U}}\| = (n + \text{tr} \mathbf{P} - 2\text{tr} \mathbf{P}^{1/2})^{1/2}$ computed by (2.11). As we can see, all three PD-based algorithms achieve a ratio of 1, while the QR-based algorithm has a ratio between 1.2 and 1.4, agreeing with the bound (3.26). Not surprisingly, the plain SVD orthonormalization yields a huge ratio of order 10^6 (not plotted) because it produces corrections of order 1 even when $\|\hat{\mathbf{U}}^T \hat{\mathbf{U}} - \mathbf{I}\|$ is only of order 10^{-6} .




Figure 5-3: The orthogonality errors after re-orthonormalization by various algorithms (left) and the ratio of the correction $\|\tilde{\mathbf{U}} - \hat{\mathbf{U}}\|$ to the minimum $\|\tilde{\mathbf{U}}_* - \hat{\mathbf{U}}\|$ computed by (2.11) (right). The ratios for SVD-EVD are of order 10^6 and thus not plotted. Here $\Delta t = 0.05$ and $n = 10$ are used for all the simulations.

² For example, Subramani (2018) uses the leapfrog scheme and Dutt (2018) uses the implicit-explicit backward difference schemes for the DO equations.

Chapter 6

Application to Ensemble Square Root Filters for Data Assimilation

Another application of the minimum-correction second-moment matching is in the ensemble square root filters for data assimilation (Evensen, 1994; Lermusiaux and Robinson, 1999; Bishop et al., 2001; Anderson, 2001; Evensen, 2003; Ott et al., 2004; Reich and Cotter, 2015). Given an n -dimensional dynamical system $d_t \mathbf{u} = \mathbf{f}(\mathbf{u})$ as well as a linear observation model $\mathbf{y} = \mathbf{H}\mathbf{u} + \mathbf{e} \in \mathbb{R}^d$ with Gaussian noise $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$, the task of data assimilation is to combine the information content of observations $\hat{\mathbf{y}}$ with that of model predictions to refine the (probabilistic) estimate of the state \mathbf{u} . At each observation time t , we denote a forecast ensemble with m realizations for \mathbf{u} by $\mathbf{U}_- \in \mathcal{M}_{n \times m}$ ¹ and its sample covariance by $\mathbf{P}_- = (1/m)\tilde{\mathbf{U}}_- \tilde{\mathbf{U}}_-^T$. Here, \mathbf{U}_- is decomposed into the mean and the fluctuation as $\mathbf{U}_- = \bar{\mathbf{u}}\mathbf{1}^T + \tilde{\mathbf{U}}_-$, where $\mathbf{1}$ is a vector with all entries being 1.

The task of filtering, the most common type of data assimilation, is to obtain an analysis ensemble $\mathbf{U}_+ \in \mathcal{M}_{n \times m}$, which integrates information in \mathbf{U}_- with the observations collected at t denoted by $\hat{\mathbf{y}}$, so that \mathbf{U}_+ better captures the underlying true state than \mathbf{U}_- can. A popular family of filters is that of ensemble square root filters, which update the mean in the same ways as a Kalman

¹ Here the role of rows and columns flips compared to previous sections, since we want to follow the notation convention in the field of data assimilation.

filter:

$$\bar{\mathbf{u}}_+ = \bar{\mathbf{u}}_- + \mathbf{K}(\hat{\mathbf{y}} - \mathbf{H}\bar{\mathbf{u}}_-), \quad \mathbf{K} = \mathbf{P}_- \mathbf{H}^T (\mathbf{H} \mathbf{P}_- \mathbf{H}^T + \mathbf{R})^{-1}, \quad (6.1)$$

and update the ensemble spread in some way such that the sample covariance of \mathbf{U}_+ matches the one given by a Kalman filter, i.e.

$$\mathbf{P}_+ = (1/m) \tilde{\mathbf{U}}_+ \tilde{\mathbf{U}}_+^T = (\mathbf{I} - \mathbf{K} \mathbf{H}) \mathbf{P}_-. \quad (6.2)$$

Many variants for obtaining such a $\tilde{\mathbf{U}}_+$ have been proposed, including the ensemble Kalman filter (Evensen, 1994, 2003), error subspace statistical estimation (Lermusiaux and Robinson, 1999; Lermusiaux, 1999), ensemble adjustment filter (Anderson, 2001) and ensemble transform filter (Bishop et al., 2001), to name a few (Vetra-Carvalho et al., 2018). Here, we demonstrate a particular variant proposed by Ott et al. (2004), which choose $\tilde{\mathbf{U}}_+$ to be the one closest to $\tilde{\mathbf{U}}_-$, reducing the task to minimum-correction second-moment matching. This choice may be seen as retaining as much physical information from the prior ensemble \mathbf{U}_- as possible (Ott et al., 2004). Here we complement it with a more quantitative motivation. Since in an ensemble square root filter only the Gaussian part of the statistics is employed for the Bayesian update, there is no information or principle that can determine how the non-Gaussian part of the statistics should be modified. Hence, in determining the posterior ensemble $\mathbf{U}_{(+)}$, any corrections made to the prior one $\mathbf{U}_{(-)}$, other than those necessary for matching the posterior mean and covariance given by the Gaussian Bayesian update, are deemed artificial and unprincipled, which we want to minimize.

More precisely, the update formula for $\tilde{\mathbf{U}}_+$ proposed by Ott et al. (2004) is

$$\tilde{\mathbf{U}}_+ = \arg \min_{\mathbf{V}: \mathbf{V} \mathbf{V}^T = m \mathbf{P}_+} \|\mathbf{V} - \tilde{\mathbf{U}}_-\|_{\mathbf{F}, \mathbf{P}_-^{-1}}. \quad (6.3)$$

corollary 2.2 implies that $\tilde{\mathbf{U}}_+ = \mathbf{A} \tilde{\mathbf{U}}_-$ with (here $\mathbf{\Gamma}_n = \mathbf{P}_-^{-1}$)

$$\mathbf{A} = \sqrt{\mathbf{\Gamma}_n}^{-1} \mathbf{A}_* \sqrt{\mathbf{\Gamma}_n} = \sqrt{\mathbf{P}_-}^{-T} (\sqrt{\mathbf{P}_-}^{-T} \mathbf{P}_+ \sqrt{\mathbf{P}_-}^{-1})^{1/2} \sqrt{\mathbf{P}_-}^{-T}. \quad (6.4)$$

This is almost the same as $\mathbf{A} = \mathbf{P}_-^{1/2} (\mathbf{P}_-^{-1/2} \mathbf{P}_+ \mathbf{P}_-^{-1/2})^{1/2} \mathbf{P}_-^{-1/2}$, which is eq. A14 in Ott et al. (2004), except that all the square root $\sqrt{\cdot}$ were unnecessarily restricted to the unique SPD one $(\cdot)^{1/2}$ in the expression used by Ott et al. (2004).

We will again use the Lorenz-96 system (5.2) as an example with exactly the same numerical setups as those in chapter 5, except that $F = 4$ is used now. Hence, we have $n = 40$, $m = 1000$. and note that n now denotes the system's dimension while m the number of realizations, differing from chapter 5. The observation data include every other variable, i.e. x_1, x_3, \dots, x_{39} , at time $t = 10, 15, 20, \dots, 50$. The observation noise covariance $\mathbf{R} = 0.01\mathbf{I}$ is assumed known.

Here four filters are tested and compared: “EnSQR” is the one proposed by Ott et al. (2004) and based on (6.3); “EnKF” is the classical ensemble Kalman filter (Evensen, 2003); “KF” is the same as EnKF except that the filtering update is replaced by fitting a Gaussian to the prior ensemble, applying the classical Kalman update and re-sampling from the posterior Gaussian. These three filters are based on brute force Monte-Carlo simulations of the Lorenz-96 system without any dimension reduction. “EnSQR-DO” is EnSQR applied to a DO simulation with a 10-dimensional reduced subspace. The performance of these four filters are compared in fig. 6-1.




Figure 6-1: Time evolution of RMSE of the various filters (left) and the ratio of the filtering correction $\|\tilde{\mathbf{U}}_+ - \tilde{\mathbf{U}}_-\|$ to the minimum $\|\tilde{\mathbf{U}}_{+*} - \tilde{\mathbf{U}}_-\|$ computed by (2.17) (right). The first three cases are without dimension reduction while the last one is reduced to a 10-dimensional subspace using DO. Here $\Delta t = 0.05$ is used for all the simulations.

The RMSE (root mean squared errors) on the left shows that all filters perform similarly except that EnSQR-DO has a larger error due to dimension reduction (the dimension of the attractor is larger than 10). The ratios of the ensemble correction (fluctuation part) $\|\tilde{\mathbf{U}}_+ - \tilde{\mathbf{U}}_-\|$ at each filtering step to the theoretical minimum $\|\tilde{\mathbf{U}}_{+*} - \tilde{\mathbf{U}}_-\|$ for posterior covariance matching are plotted on the right. The minimum is computed by (2.17) being $\|\tilde{\mathbf{U}}_{+*} - \tilde{\mathbf{U}}_-\|^2 = \text{tr}(\mathbf{I} + \tilde{\mathbf{P}}_\Gamma - 2\tilde{\mathbf{P}}_\Gamma^{1/2})$, where

$\tilde{\mathbf{P}}_{\Gamma} = \sqrt{\mathbf{P}_{-}}^{-\text{T}} \mathbf{P}_{+} \sqrt{\mathbf{P}_{-}}^{-1}$. As we can see, the two minimum-correction-based filters (EnSQR, EnSQR-DO) achieve the unit ratio while the other two do not. KF has the largest ratio because its re-sampling step completely breaks any realization correspondence between the prior and posterior ensemble. In contrast, EnKF adopts a realization-wise update strategy, which yields a ratio between those of KF and of the minimum-correction filters.

Chapter 7

Conclusions

In this paper, we develop the theory for minimum-correction second-moment matching. We solve the optimization analytically for both the full-rank and rank-deficient case. The finite-dimensional result is generalized to the infinite-dimensional setting and a connection to optimal transport is drawn. We also show how the general problem can reduce to the familiar one of optimal orthonormalization and we accordingly generalize the polar decomposition, which is known to solve this special case.

Besides the theory, we devise the computation strategies for numerically computing the optimizer. We show the instrumental role played by the algorithms for polar decomposition and SPD matrix square root and we analyze existing ones in the literature. We modify two Newton iterations deemed unstable before to significantly improve their stability and the two new multi-purpose schemes (3.19) and (3.20) can be used to efficiently compute both the orthogonal polar factor and the SPD square root. These iterations play a key role in minimum-correction second-moment matching using algorithm \mathbf{A}_* -NtSqr and $\tilde{\mathbf{U}}_*$ -NtPD in table 3.1 and in their counterparts PD-NtSqr, PD-NtISqr, and PD-NtPD in table 4.1 for orthonormalization. We verify the higher performance of the new algorithms using benchmarks with random matrices.

Finally, we showcase results in two applications. In reduced subspace tracking for uncertainty quantification, we maintain the numerical orthogonality and continuity of a time-varying orthonormal basis. In an ensemble square root filter for data assimilation, we obtain the unique posterior

ensemble that matches a given covariance matrix while also minimizing its distance to the prior one.

Appendix A

Computing the Matrix Gradients

In (C.1) and (2.3), we need to compute the matrix gradient of a scalar function analytically. In general, given a scalar function $f : \mathcal{M}_{n \times m}(\mathbb{R}) \rightarrow \mathbb{R}$ that maps an n -by- m matrix \mathbf{A} to a real number $f(\mathbf{A})$, the matrix gradient $\nabla_{\mathbf{A}} f$ is defined as the n -by- m matrix whose (i, j) -th element is $\partial f / \partial A_{ij}$. Therefore, the matrix gradient is simply the particular reshaping of the normal gradient vector that matches the variable layout in \mathbf{A} . This reshaping allows us to take into account the special structures of the matrix operations involved in $f(\cdot)$, such as matrix multiplication, taking trace and determinant, etc., to express the gradient analytically in a compact and elegant form.

The analytic expressions of many typical matrix gradients of a scalar function encountered in practice have been listed in (Petersen and Pedersen, 2012, ch. 2). Here we show how we use the results provided there to derive (C.1) and (2.3).

In (C.1), we need $\nabla_{\mathbf{A}} \text{tr}(\mathbf{A}^T \mathbf{P} \mathbf{A})$ and $\nabla_{\mathbf{A}} \text{tr}(\mathbf{A}^T \mathbf{P} \mathbf{A} \mathbf{A}^T \mathbf{P} \mathbf{A})$. Eq. (111) in (Petersen and Pedersen, 2012, sec. 2.5.2) gives $\nabla_{\mathbf{X}} \text{tr}(\mathbf{X} \mathbf{B} \mathbf{X}^T) = \mathbf{X}(\mathbf{B}^T + \mathbf{B})$. Since \mathbf{P} is symmetric, we have

$$\nabla_{\mathbf{A}} \text{tr}(\mathbf{A}^T \mathbf{P} \mathbf{A}) = [\nabla_{\mathbf{A}^T} \text{tr}(\mathbf{A}^T \mathbf{P} \mathbf{A})]^T = (2\mathbf{A}^T \mathbf{P})^T = 2\mathbf{P} \mathbf{A}. \quad (\text{A.1})$$

Eq. (123) in (Petersen and Pedersen, 2012, sec. 2.5.3) gives

$$\begin{aligned}\nabla_{\mathbf{X}}\text{tr}(\mathbf{B}^T\mathbf{X}^T\mathbf{C}\mathbf{X}\mathbf{X}^T\mathbf{C}\mathbf{X}\mathbf{B}) &= \mathbf{C}\mathbf{X}\mathbf{X}^T\mathbf{C}\mathbf{X}\mathbf{B}\mathbf{B}^T + \mathbf{C}^T\mathbf{X}\mathbf{B}\mathbf{B}^T\mathbf{X}^T\mathbf{C}^T\mathbf{X} \\ &\quad + \mathbf{C}\mathbf{X}\mathbf{B}\mathbf{B}^T\mathbf{X}^T\mathbf{C}\mathbf{X} + \mathbf{C}^T\mathbf{X}\mathbf{X}^T\mathbf{C}^T\mathbf{X}\mathbf{B}\mathbf{B}^T.\end{aligned}$$

Since $\mathbf{B} = \mathbf{I}$ and \mathbf{P} is symmetric here, we have

$$\nabla_{\mathbf{A}}\text{tr}(\mathbf{A}^T\mathbf{P}\mathbf{A}\mathbf{A}^T\mathbf{P}\mathbf{A}) = 4\mathbf{P}\mathbf{A}\mathbf{A}^T\mathbf{P}\mathbf{A}. \quad (\text{A.2})$$

(A.1) and (A.2) together yield (C.1).

In (2.3), we need $\nabla_{\mathbf{A}}\text{tr}(\mathbf{P}\mathbf{A})$, $\nabla_{\mathbf{A}}\text{tr}(\mathbf{\Lambda}^T\mathbf{A}^T\mathbf{P}\mathbf{A})$ and $\nabla_{\mathbf{\Lambda}}\text{tr}(\mathbf{\Lambda}^T(\mathbf{A}^T\mathbf{P}\mathbf{A} - \mathbf{I}))$. The first two can be obtained by eq. (103) in (Petersen and Pedersen, 2012, sec. 2.5.1): $\nabla_{\mathbf{X}}\text{tr}(\mathbf{X}^T\mathbf{A}) = \mathbf{A}$. Therefore,

$$\nabla_{\mathbf{\Lambda}}\text{tr}(\mathbf{\Lambda}^T(\mathbf{A}^T\mathbf{P}\mathbf{A} - \mathbf{I})) = \mathbf{A}^T\mathbf{P}\mathbf{A} - \mathbf{I}. \quad (\text{A.3})$$

Also, since \mathbf{P} is symmetric, we have

$$\nabla_{\mathbf{A}}\text{tr}(\mathbf{P}\mathbf{A}) = \nabla_{\mathbf{A}}\text{tr}(\mathbf{A}^T\mathbf{P}) = \mathbf{P}. \quad (\text{A.4})$$

Eq. (117) in (Petersen and Pedersen, 2012, sec. 2.5.2) gives $\nabla_{\mathbf{X}}\text{tr}(\mathbf{X}^T\mathbf{B}\mathbf{X}\mathbf{C}) = \mathbf{B}\mathbf{X}\mathbf{C} + \mathbf{B}^T\mathbf{X}\mathbf{C}^T$. Since both $\mathbf{\Lambda}$ and \mathbf{P} are symmetric, we have

$$\nabla_{\mathbf{A}}\text{tr}(\mathbf{\Lambda}^T\mathbf{A}^T\mathbf{P}\mathbf{A}) = \nabla_{\mathbf{A}}\text{tr}(\mathbf{A}^T\mathbf{P}\mathbf{A}\mathbf{\Lambda}) = 2\mathbf{P}\mathbf{A}\mathbf{\Lambda}. \quad (\text{A.5})$$

(A.3), (A.4) and (A.5) together yield (2.3).

Appendix B

Successive First Order Approximation for Matrix Square Root

Here we show a motivation of the iteration (3.20) as a successive first order approximation to the solution to $\mathbf{G}(\mathbf{A}) = \mathbf{A}^T \mathbf{P} \mathbf{A} - \mathbf{I} = \mathbf{0}$.

Since we want $\mathbf{A}^T \mathbf{P} \mathbf{A} = \mathbf{I}$, we have $\mathbf{A}^T \mathbf{A} + \mathbf{A}^T (\mathbf{P} - \mathbf{I}) \mathbf{A} = \mathbf{I}$ and thus

$$(\mathbf{A}^T - \mathbf{I}) + (\mathbf{A} - \mathbf{I}) + (\mathbf{A}^T - \mathbf{I})(\mathbf{A} - \mathbf{I}) + \mathbf{A}^T (\mathbf{P} - \mathbf{I}) \mathbf{A} = \mathbf{0}.$$

If we assume both \mathbf{A} and \mathbf{P} are close to \mathbf{I} , we have the following first-order approximation to the above: $(\mathbf{A}^T - \mathbf{I}) + (\mathbf{A} - \mathbf{I}) = -(\mathbf{P} - \mathbf{I})$. Therefore, for \mathbf{P} close to \mathbf{I} , \mathbf{A} can be as close to \mathbf{I} as $\|\mathbf{A} - \mathbf{I}\| \sim (1/2)\|\mathbf{P} - \mathbf{I}\|$. Since we are seeking the SPD square root of \mathbf{P}^{-1} , symmetry can be imposed on \mathbf{A} , which gives

$$\mathbf{A} = \mathbf{I} - (1/2)(\mathbf{P} - \mathbf{I}). \tag{B.1}$$

Note that this \mathbf{A} is only a first-order approximation of the optimizer, so $\mathbf{A}^T \mathbf{P} \mathbf{A} - \mathbf{I} = 0$ will not

be satisfied up to round-off errors. In fact, we can check

$$\begin{aligned}
\mathbf{A}^T \mathbf{P} \mathbf{A} - \mathbf{I} &= [\mathbf{I} - (1/2)(\mathbf{P} - \mathbf{I})]^T \mathbf{P} [\mathbf{I} - (1/2)(\mathbf{P} - \mathbf{I})] - \mathbf{I} \\
&= \mathbf{P} - (\mathbf{P} - \mathbf{I})\mathbf{P} + (1/4)(\mathbf{P} - \mathbf{I})\mathbf{P}(\mathbf{P} - \mathbf{I}) - \mathbf{I} \\
&= -(3/4)(\mathbf{P} - \mathbf{I})^2 + (1/4)(\mathbf{P} - \mathbf{I})^3.
\end{aligned}$$

Therefore, although $\mathbf{U}\mathbf{A}$ is not exactly orthogonal, its deviation from orthogonality gets squared and diminished: $\|\mathbf{A}^T \mathbf{P} \mathbf{A} - \mathbf{I}\| \sim (3/4) \|\mathbf{P} - \mathbf{I}\|^2$. Therefore, if we treat $\mathbf{A}^T \mathbf{P} \mathbf{A}$ as a new “ \mathbf{P} ” and perform the same first-order approximation to improve the solution, we obtain the following iteration (with $\mathbf{P}_0 = \mathbf{P}$):

$$\mathbf{A}_{k+1} = \mathbf{I} - (1/2)(\mathbf{P}_k - \mathbf{I}), \quad \mathbf{P}_{k+1} = \mathbf{A}_{k+1}^T \mathbf{P}_k \mathbf{A}_{k+1}.$$

If we introduce $\mathbf{X}_{k+1} = \mathbf{X}_k \mathbf{A}_{k+1}$ (with $\mathbf{X}_0 = \mathbf{I}$) as the composition of previous corrections, we have $\mathbf{X}_{k+1} = \mathbf{X}_k - (1/2)\mathbf{X}_k(\mathbf{X}_k^T \mathbf{P} \mathbf{X}_k - \mathbf{I})$, which recovers (3.20).

Appendix C

Gradient Descent for Matrix Square Root

Feppon and Lermusiaux (2018b) proposes to solve the equation $\mathbf{X}^\top \mathbf{P} \mathbf{X} = \mathbf{I}$ for an \mathbf{X} close to \mathbf{I} by solving the unconstrained optimization $\arg \min_{\mathbf{X}} \|\mathbf{X}^\top \mathbf{P} \mathbf{X} - \mathbf{I}\|_{\mathbb{F}}^2$ iteratively with the initial guess $\mathbf{X}_0 = \mathbf{I}$.

The objective function attains its global minimum 0 at every \mathbf{X} solves $\mathbf{X}^\top \mathbf{P} \mathbf{X} = \mathbf{I}$. The hope is that if we use an iterative scheme and start from \mathbf{I} , the iterates will converge to a solution \mathbf{X} close to \mathbf{I} . Feppon and Lermusiaux (2018b) proposes to use a simple gradient descent iteration. Since $f(\mathbf{X}) = \|\mathbf{X}^\top \mathbf{P} \mathbf{X} - \mathbf{I}\|^2 = \text{tr}(\mathbf{X}^\top \mathbf{P} \mathbf{X} \mathbf{X}^\top \mathbf{P} \mathbf{X}) - 2\text{tr}(\mathbf{X}^\top \mathbf{P} \mathbf{X}) + \text{tr}(\mathbf{I})$, we can compute the gradient analytically (see appendix A) as

$$\nabla_{\mathbf{X}} f = 4\mathbf{P} \mathbf{X} (\mathbf{X}^\top \mathbf{P} \mathbf{X} - \mathbf{I}). \quad (\text{C.1})$$

Therefore, the gradient descent iteration takes the form

$$\mathbf{X}_{n+1} = \mathbf{X}_n - \gamma \mathbf{P} \mathbf{X}_n (\mathbf{X}_n^\top \mathbf{P} \mathbf{X}_n - \mathbf{I}) \quad (\text{C.2})$$

with γ being a tunable step size. This is Algorithm 3 in (Feppon and Lermusiaux, 2018b, sec. 3.5).

However, Feppon and Lermusiaux (2018b) does not mention how to pick γ , when the iteration converges and how fast.

We can analyze this iteration by exactly the same techniques used in section 3.2. First, we have $\mathbf{X}_k \rightarrow \mathbf{P}^{-1/2}\mathbf{Q}_0$ for

$$\mathbf{X}_0 = \mathbf{S}_0\mathbf{Q}_0, \quad \mathbf{S}_0 \in \mathcal{C}_P, \quad \|\mathbf{X}_0\|_2 \leq 1/\|\mathbf{P}\|_2^{1/2}, \quad \gamma \leq 1/(2\|\mathbf{P}\|_2), \quad (\text{C.3})$$

since the matrix iteration can be reduced to the scalar one $\sigma_{k+1} = \sigma_k - \gamma\lambda\sigma_k(\sigma_k^2 - 1)^1$, where λ is an eigenvalue of \mathbf{P} and $\sigma_k/\sqrt{\lambda}$ is the corresponding singular value of \mathbf{X}_k . Next, the stability analysis shows that the Fréchet derivative at the limit \mathbf{X}_* is

$$\mathbf{D}\mathbf{F}(\mathbf{X}) = \mathbf{X} - \gamma\mathbf{\Lambda}\mathbf{X} - \gamma\mathbf{\Lambda}^{1/2}\mathbf{X}^T\mathbf{\Lambda}^{1/2}, \quad (\text{C.4})$$

which has eigenvectors and eigenvalues as

$$\begin{aligned} \mathbf{Z}_{ij} &= \mathbf{E}_{ij} + \sqrt{\lambda_j/\lambda_i}\mathbf{E}_{ji}, \quad \mu_{ij} = 1 - \gamma(\lambda_i + \lambda_j), \quad i \leq j, \\ \mathbf{Z}_{ij} &= \mathbf{E}_{ij} - \sqrt{\lambda_i/\lambda_j}\mathbf{E}_{ji}, \quad \mu_{ij} = 1, \quad i > j. \end{aligned} \quad (\text{C.5})$$

Here the λ_i 's are the diagonal entries of $\mathbf{\Lambda}$ and are also the eigenvalues of \mathbf{P} . Therefore, for $\rho(\mathbf{D}\mathbf{F}) < 1$, we need $\gamma < 1/\|\mathbf{P}\|_2$, which is less stringent than (C.3). Moreover, the iteration converges linearly with the asymptotic error diminishing coefficient not smaller than $(\lambda_{\max} - \lambda_{\min})/(\lambda_{\max} + \lambda_{\min})$. Therefore, the gradient descent is considered unattractive here compared to the quadratically converging iterations (3.19) and (3.20).

¹ Since $(1 - \sigma_{k+1}) = (1 - \sigma_k)(1 - \gamma\lambda\sigma_k(\sigma_k^2 + 1))$, if $0 < \sigma_0 \leq 1$ and $0 < \gamma \leq 1/(2\lambda)$, σ_k monotonically increases and converges to 1.

Bibliography

- Absil, P.-A., Mahony, R., and Sepulchre, R. (2009). *Optimization Algorithms on Matrix Manifolds*. Princeton University Press.
- Absil, P.-A. and Malick, J. (2012). Projection-Like Retractions on Matrix Manifolds. *SIAM Journal on Optimization*, 22(1):135–158.
- Anderson, J. L. (2001). An Ensemble Adjustment Kalman Filter for Data Assimilation. *Monthly weather review*, 129(12):2884–2903.
- Benamou, J.-D. and Brenier, Y. (2000). A Computational Fluid Mechanics Solution to the Monge-Kantorovich Mass Transfer Problem. *Numerische Mathematik*, 84(3):375–393.
- Bergemann, K., Gottwald, G., and Reich, S. (2009). Ensemble propagation and continuous matrix factorization algorithms. *Quarterly Journal of the Royal Meteorological Society*, 135(643):1560–1572.
- Bishop, C. H., Etherton, B. J., and Majumdar, S. J. (2001). Adaptive Sampling with the Ensemble Transform Kalman Filter. Part I: Theoretical Aspects. *Monthly weather review*, 129(3):420–436.
- Brenier, Y. (1991). Polar Factorization and Monotone Rearrangement of Vector-Valued Functions. *Communications on Pure and Applied Mathematics*, 44(4):375–417.
- Brockett, R. W. (1991). Dynamical Systems that Learn Subspaces. In *Mathematical System Theory*, pages 579–592. Springer.
- Chandrasekaran, S. and Ipsen, I. C. (1994). Backward Errors for Eigenvalue and Singular Value Decompositions. *Numerische Mathematik*, 68(2):215–223.
- Cheng, M., Hou, T. Y., and Zhang, Z. (2013). A Dynamically Bi-Orthogonal Method for Time-Dependent Stochastic Partial Differential Equations I: Derivation and Algorithms. *Journal of Computational Physics*, 242:843–868.

- Dieci, L., Russell, R. D., and Van Vleck, E. S. (1994). Unitary Integrators and Applications to Continuous Orthonormalization Techniques. *SIAM Journal on Numerical Analysis*, 31(1):261–281.
- Dowson, D. and Landau, B. (1982). The Fréchet Distance between Multivariate Normal Distributions. *Journal of Multivariate Analysis*, 12(3):450–455.
- Dutt, A. (2018). High Order Stochastic Transport and Lagrangian Data Assimilation. Master’s thesis, Massachusetts Institute of Technology, Department of Mechanical Engineering, Cambridge, Massachusetts.
- Edelman, A., Arias, T. A., and Smith, S. T. (1998). The Geometry of Algorithms with Orthogonality Constraints. *SIAM journal on Matrix Analysis and Applications*, 20(2):303–353.
- Evensen, G. (1994). Sequential Data Assimilation with a Nonlinear Quasi-Geostrophic Model Using Monte Carlo Methods to Forecast Error Statistics. *Journal of Geophysical Research: Oceans*, 99(C5):10143–10162.
- Evensen, G. (2003). The Ensemble Kalman Filter: Theoretical Formulation and Practical Implementation. *Ocean dynamics*, 53(4):343–367.
- Feppon, F. and Lermusiaux, P. F. J. (2018a). A Geometric Approach to Dynamical Model-Order Reduction. *SIAM Journal on Matrix Analysis and Applications*, 39(1):510–538.
- Feppon, F. and Lermusiaux, P. F. J. (2018b). Dynamically Orthogonal Numerical Schemes for Efficient Stochastic Advection and Lagrangian Transport. *SIAM Review*, 60(3):595–625.
- Feppon, F. and Lermusiaux, P. F. J. (2019). The Extrinsic Geometry of Dynamical Systems Tracking Nonlinear Matrix Projections. *SIAM Journal on Matrix Analysis and Applications*, 40(2):814–844.
- Golub, G. H. and Van Loan, C. F. (2013). *Matrix Computations*. JHU Press, 4th edition.
- Govaerts, W. and Werner, B. (1995). Continuous Bordering of Matrices and Continuous Matrix Decompositions. *Numerische Mathematik*, 70(3):303–310.
- Hairer, E., Lubich, C., and Wanner, G. (2006). *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*, volume 31. Springer Science & Business Media.
- Higham, D. J. (1996). Runge-Kutta Type Methods for Orthogonal Integration. *Applied Numerical Mathematics*, 22(1-3):217–223.
- Higham, D. J. (1997). Time-Stepping and Preserving Orthonormality. *BIT Numerical Mathematics*, 37(1):24–36.

- Higham, N. J. (1986a). Computing the Polar Decomposition with Applications. *SIAM Journal on Scientific and Statistical Computing*, 7(4):1160–1174.
- Higham, N. J. (1986b). Newton’s Method for the Matrix Square Root. *Mathematics of Computation*, 46(174):537–549.
- Higham, N. J. (1988). Matrix Nearness Problems and Applications. Technical report, Department of Mathematics, University of Manchester.
- Higham, N. J., Mehl, C., and Tisseur, F. (2010). The Canonical Generalized Polar Decomposition. *SIAM Journal on Matrix Analysis and Applications*, 31(4):2163–2180.
- Higham, N. J. and Schreiber, R. S. (1990). Fast Polar Decomposition of an Arbitrary Matrix. *SIAM Journal on Scientific and Statistical Computing*, 11(4):648–655.
- Koch, O. and Lubich, C. (2007). Dynamical Low-Rank Approximation. *SIAM Journal on Matrix Analysis and Applications*, 29(2):434–454.
- Laub, A. J. (2005). *Matrix Analysis for Scientists and Engineers*. SIAM.
- Lermusiaux, P. F. J. (1997). *Error Subspace Data Assimilation Methods for Ocean Field Estimation: Theory, Validation and Applications*. Harvard University.
- Lermusiaux, P. F. J. (1999). Data Assimilation via Error Subspace Statistical Estimation, Part II: Mid-Atlantic Bight Shelfbreak Front Simulations, and ESSE Validation. *Monthly Weather Review*, 127(7):1408–1432.
- Lermusiaux, P. F. J. (2001). Evolving the Subspace of the Three-Dimensional Multiscale Ocean Variability: Massachusetts Bay. *Journal of Marine Systems*, 29(1):385–422.
- Lermusiaux, P. F. J. and Robinson, A. R. (1999). Data Assimilation via Error Subspace Statistical Estimation, Part I: Theory and Schemes. *Monthly Weather Review*, 127(7):1385–1407.
- Lorenz, E. N. (1996). Predictability: A Problem Partly Solved. In *Proc. Seminar on Predictability*, volume 1.
- Musharbash, E. and Nobile, F. (2018). Dual Dynamically Orthogonal Approximation of Incompressible Navier Stokes Equations with Random Boundary Conditions. *Journal of Computational Physics*, 354:135–162.
- Olkin, I. and Pukelsheim, F. (1982). The Distance between Two Random Vectors with Given Dispersion Matrices. *Linear Algebra and its Applications*, 48:257–263.

- Ott, E., Hunt, B. R., Szunyogh, I., Zimin, A. V., Kostelich, E. J., Corazza, M., Kalnay, E., Patil, D., and Yorke, J. A. (2004). A Local Ensemble Kalman Filter for Atmospheric Data Assimilation. *Tellus A*, 56(5):415–428.
- Petersen, K. B. and Pedersen, M. S. (2012). *The Matrix Cookbook*. Technical University of Denmark.
- Philippe, B. (1987). An Algorithm to Improve Nearly Orthonormal Sets of Vectors on a Vector Processor. *SIAM Journal on Algebraic Discrete Methods*, 8(3):396–403.
- Reich, S. and Cotter, C. (2015). *Probabilistic Forecasting and Bayesian Data Assimilation*. Cambridge University Press.
- Rheinboldt, W. C. (1988). On the Computation of Multi-Dimensional Solution Manifolds of Parametrized Equations. *Numerische Mathematik*, 53(1-2):165–181.
- Sapsis, T. P. and Lermusiaux, P. F. J. (2009). Dynamically Orthogonal Field Equations for Continuous Stochastic Dynamical Systems. *Physica D: Nonlinear Phenomena*, 238(23–24):2347–2360.
- Schatzman, M. (2002). *Numerical Analysis: A Mathematical Introduction*. Oxford University Press.
- Strogatz, S. H. (1994). *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*. Westview Press.
- Subramani, D. and Lermusiaux, P. F. J. (2019). Probabilistic Ocean Predictions with Dynamically-Orthogonal Primitive Equations. *In preparation*.
- Subramani, D. N. (2018). *Probabilistic Regional Ocean Predictions: Stochastic Fields and Optimal Planning*. PhD thesis, Massachusetts Institute of Technology, Department of Mechanical Engineering, Cambridge, Massachusetts.
- Tagare, H. D. (2011). Notes on Optimization on Stiefel Manifolds. Yale University.
- Trefethen, L. N. and Bau III, D. (1997). *Numerical Linear Algebra*, volume 50. SIAM.
- Ueckermann, M., Lermusiaux, P., and Sapsis, T. (2011). Numerical Schemes and Computational Studies for Dynamically Orthogonal Equations. MSEAS Report 11, Department of Mechanical Engineering, Massachusetts Institute of Technology.
- Ueckermann, M. P., Lermusiaux, P. F. J., and Sapsis, T. P. (2013). Numerical Schemes for Dynamically Orthogonal Equations of Stochastic Fluid and Ocean Flows. *Journal of Computational Physics*, 233:272–294.

- Vetra-Carvalho, S., van Leeuwen, P. J., Nerger, L., Barth, A., Altaf, M. U., Brasseur, P., Kirchgessner, P., and Beckers, J.-M. (2018). State-of-the-Art Stochastic Data Assimilation Methods for High-Dimensional Non-Gaussian Problems. *Tellus A: Dynamic Meteorology and Oceanography*, 70(1):1–43.
- Villani, C. (2008). *Optimal Transport: Old and New*, volume 338. Springer Science & Business Media.
- Wen, Z. and Yin, W. (2013). A Feasible Method for Optimization with Orthogonality Constraints. *Mathematical Programming*, 142(1-2):397–434.
- Yang, H. and Li, H. (2008). Weighted Polar Decomposition and WGL Partial Ordering of Rectangular Complex Matrices. *SIAM Journal on Matrix Analysis and Applications*, 30(2):898–924.