



HAL
open science

Low-rank plus diagonal approximations for Riccati-like matrix differential equations

Silvère Bonnabel, Marc Lambert, Francis Bach

► **To cite this version:**

Silvère Bonnabel, Marc Lambert, Francis Bach. Low-rank plus diagonal approximations for Riccati-like matrix differential equations. *SIAM Journal on Matrix Analysis and Applications*, 2024, 45 (3), pp.1669-1688. <10.1137/23M1587610>. <hal-04621158>

HAL Id: hal-04621158

<https://inria.hal.science/hal-04621158v1>

Submitted on 28 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Low-rank plus diagonal approximations for Riccati-like matrix differential equations

Silvère Bonnabel

MINES ParisTech, PSL University, Center for robotics
`silvere.bonnabel@mines-paristech.fr`

Marc Lambert

INRIA - Ecole Normale Supérieure - PSL Research university
DGA/CATOD, Centre d'Analyse Technico-Opérationnelle de Défense
`marc.lambert@inria.fr`

Francis Bach

INRIA - Ecole Normale Supérieure - PSL Research university
`francis.bach@inria.fr`

Abstract

We consider the problem of computing tractable approximations of time-dependent $d \times d$ large positive semi-definite (PSD) matrices defined as solutions of a matrix differential equation. We propose to use “low-rank plus diagonal” PSD matrices as approximations that can be stored with a memory cost being linear in the high dimension d . To constrain the solution of the differential equation to remain in that subset, we project the derivative at all times onto the tangent space to the subset, following the methodology of dynamical low-rank approximation. We derive a closed-form formula for the projection, and show that after some manipulations it can be computed with a numerical cost being linear in d , allowing for tractable implementation. Contrary to previous approaches based on pure low-rank approximations, the addition of the diagonal term allows for our approximations to be invertible matrices, that can moreover be inverted with linear cost in d . We apply the technique to Riccati-like equations, then to two particular problems. Firstly a low-rank approximation to our recent Wasserstein gradient flow for Gaussian approximation of posterior distributions in approximate Bayesian inference, and secondly a novel low-rank approximation of the Kalman filter for high-dimensional systems. Numerical simulations illustrate the results.

1 Introduction

Positive semi-definite (PSD) matrices $X \in \mathbb{R}^{d \times d}$ have storage cost in d^2 , and the computation cost associated with typical matrix operations is in d^3 . Those two aspects may come as a limitation in various applications where the dimension d is very large. A sensible approach is thus to work with low-rank approximations instead. A rank p approximation of a PSD matrix X can be factored as $Y = URU^T$ where $U \in \mathbb{R}^{d \times p}$ has orthonormal columns and R is a $p \times p$ positive definite matrix of much-reduced size, letting $p \ll d$. We denote by $S^+(p, d)$ the set of rank- p PSD matrices. A somewhat richer set, that will be the object of this paper, is the set of “low-rank plus diagonal” matrices, that we denote by $S_{\text{diag}}^+(p, d)$, composed of matrices of the form $Y = URU^T + \psi$ with ψ a diagonal matrix with strictly positive diagonal elements, making Y invertible. We then speak of FA decomposition, in relation to the problem of factor analysis (FA) [13]. In the case where the

diagonal matrix ψ is made isotropic, that is, is taken of the form $\psi = sI_d$, $s > 0$, we speak of PPCA decomposition, in relation to the problem of probabilistic principal component analysis (PPCA) [29], and denote this smaller set by $S_{\text{isot}}^+(\text{p}, \text{d})$. A well-known "static" approximation problem consists in computing the closest approximating matrix Y to some large PSD matrix X in the sense of the Frobenius norm, that is, solving

$$\min_{Y \in \mathcal{M}} \|Y - X\|, \quad (1)$$

where \mathcal{M} may denote $S^+(\text{p}, \text{d})$, $S_{\text{diag}}^+(\text{p}, \text{d})$, or $S_{\text{isot}}^+(\text{p}, \text{d})$. If $\mathcal{M} = S^+(\text{p}, \text{d})$, the problem is solved by a singular value decomposition where eigenvalues are truncated after the p -th one, a classical result of matrix analysis known as the Eckart–Young–Mirsky theorem [10], see also [15]. If $\mathcal{M} = S_{\text{diag}}^+(\text{p}, \text{d})$, we recover the problem known as "minimum residual factor analysis", see [27, 3]. The problem is not solvable in closed form, but numerical methods have been developed, see e.g. [9].

In this paper, however, we consider a different problem. We consider positive semi-definite matrices $X(t) \in \mathbb{R}^{d \times d}$ defined as solutions of matrix differential equations in high dimension. Instead of attacking (1) directly, we try to solve at all times

$$\min_{\dot{Y}(t) \in \mathcal{T}_{Y(t)}\mathcal{M}} \|\dot{Y}(t) - \dot{X}(t)\|. \quad (2)$$

Technically, the problem differs from (1) as the derivatives live in a different space, namely the tangent spaces. The rationale is as follows. To handle possibly high dimension d , we would like to maintain a storage cost being linear in d —which is achieved by letting $Y(t) \in \mathcal{M}$ at all times—and a computation cost being also linear in d when incrementally evolving our approximation $Y(t)$, hence the need to solve (2) efficiently. The solution to

$$\dot{X}(t) = F(X(t)) \quad (3)$$

may then be approximated by letting at all times $\dot{Y}(t)$ be the solution to (2), where $\dot{X}(t) = F(X(t))$ is replaced by $F(Y(t))$. Indeed, even if $Y(t) \in \mathcal{M}$, the matrix $F(Y(t))$ generally points to a direction that makes $Y(t)$ step out of \mathcal{M} , hence the need to project it onto the tangent space $\mathcal{T}_{Y(t)}\mathcal{M}$. Although $\mathcal{T}_{Y(t)}\mathcal{M}$ is a vector space for each $Y(t)$, solving (2) exactly while maintaining operations being linear in d is not straightforward, and this is the object of the present paper.

Let us first consider $\mathcal{M} = S^+(\text{p}, \text{d})$. Low-rank approximations $Y(t)$ to $X(t)$ conveniently write $Y(t) = U(t)R(t)U(t)^T$ with $U(t) \in \mathbb{R}^{d \times p}$ a matrix whose columns are orthonormal and $R(t) \in S^+(\text{p}, \text{p})$ a small-size positive definite (PD) symmetric matrix, as advocated in [16, 4, 19]. [16, 19] address the following projection problem

$$\min_{\dot{Y}(t) \in \mathcal{T}_{Y(t)}S^+(\text{p}, \text{d})} \|\dot{Y}(t) - \dot{X}(t)\| \quad (4)$$

in the sense of the Frobenius norm, where $\mathcal{T}_Y S^+(\text{p}, \text{d})$ denotes the tangent space to $S^+(\text{p}, \text{d})$ at $Y \in S^+(\text{p}, \text{d})$. [16] derives expressions for $\dot{U}(t)$ and $\dot{R}(t)$ such that $\dot{Y}(t)$ corresponds to the solution of (4) indeed.

Historically, [5, 7, 19] proposed approximations to the continuous-time low-rank Riccati equation using differential geometry, which were in particular applied to approximating the Linblad equations in quantum physics. The work [7, 19] is in line with prior results of [16], where essentially one seeks to orthogonally project the tangent vector in the sense of the Euclidean norm, whereas the approach of [5] builds on the geometry of [4]. A judicious numerical time integrator was proposed in [22] with recent extensions to matrices and tensors that preserve symmetry in [8].

The present paper builds on [16, 7, 19], providing an invertible covariance matrix, by approximating derivatives on $S_{\text{diag}}^+(\mathbf{p}, d)$ and $S_{\text{isot}}^+(\mathbf{p}, d)$ instead. This has not been addressed to our knowledge, and we provide closed-form formulas for the projection of the tangent vector. Although the method is general, implementing the projection depends on the underlying differential equation being approximated. Of particular interest is the Riccati differential equation, that appears in a number of applications, e.g., Kalman filtering, linear quadratic control, in stability analysis via the Lyapunov function, or in computational statistics to estimate covariance matrices, see [17]. The interests of the ‘‘low-rank plus diagonal’’ framework in the context of matrix dynamical approximation are as follows:

- It allows for manipulation of *invertible* matrices, whereas low-rank approximations are inherently non-invertible. Moreover, using Woodbury’s lemma, inversion of such an approximation is feasible at all times, see (27) below.
- Approximation with a full rank matrix is desirable in estimation problems where Y encodes a covariance matrix. A rank-deficient covariance may lead to over-confident estimates. In Kalman filtering, this may lead the state variable estimates to spuriously drift, as shall be shown in simulations.
- It ensures more flexibility while retaining a storage cost linear in d . Notably, it can capture the individual variances of the variables for a large covariance matrix Y .
- Y admits a probabilistic interpretation: $URU^T + sI_d$ is the covariance of $x = Uz + \nu$, with $z \sim \mathcal{N}(0, R)$ a small-dimensional latent variable mapped via U into a subspace of \mathbb{R}^d , plus ν an isotropic noise of magnitude s , [29, 13].

Our contributions and the organization of the paper are as follows. Previous work on low-rank approximation of matrix differential equations [16, 19] is recapped in Section 2. Section 3 extends the results to the low-rank plus diagonal case: We solve exactly in closed form for each t the following extension of problem (4):

$$\min_{\dot{Y}(t) \in \mathcal{T}_{Y(t)} S_{\text{diag}}^+(\mathbf{p}, d)} \|\dot{Y}(t) - \dot{X}(t)\|, \quad (5)$$

for the set $S_{\text{diag}}^+(\mathbf{p}, d)$ and also for $S_{\text{isot}}^+(\mathbf{p}, d)$. In Section 4, we discuss implementation in high dimension and show the approximation can be obtained exactly with a computational cost being linear in d . In Section 5, we leverage the result to provide a tractable approximation to the Riccati equation. This is applied to two problems. First, we derive a low-rank approximation to our Wasserstein gradient flow of [18]. Then, we propose a novel low-rank plus diagonal Kalman filter and illustrate the results on a tutorial example inspired by robotics. The code is made publicly available.

As concerns low-rank techniques for Kalman filtering, we note the problem has arisen in data assimilation for weather forecasting and oceanography [11], [30], where the state is driven by a partial differential equation, and is thus encoded by a high dimensional vector (say, up to 1 million). Since the Kalman filter needs to store the covariance matrix of the estimates, it meets the computer’s memory limits when the dimension is high. Different variants have been proposed to tackle this memory problem. The SEEK filter [30] is based on an SVD decomposition of the covariance matrix and provides a low-rank Riccati update in discrete time. The ensemble Kalman filter [11, 20] approximates the Riccati equation using Monte Carlo sampling. Very recently, and closely related to the present paper, a method for efficient, approximate Gaussian (Kalman) filtering, including

smoothing and marginal likelihood computation was developed using low-rank approximations on $S^+(p, d)$, see [26]. In robotics, [28] proposed the sparse information filter (SEIF) filter to process large maps for simultaneous localization and mapping.

2 Reminders on low-rank approximation

We start with the geometry of $S^+(p, d)$, and then recap existing results on dynamical low-rank approximation.

2.1 Geometry of $S^+(p, d)$

Any matrix $Y \in S^+(p, d)$ may be written as $Y = URU^T$ where $R \in S^+(p, p)$ is a small size positive definite matrix, and $U \in V(p, d)$ where $V(p, d)$ denotes the set of $d \times p$ matrices with p orthonormal columns, hence satisfying $U^T U = I_p$, called the Stiefel manifold. Letting $\Pi_U = UU^T$ be the projector onto the span of U and $\Pi_U^\perp = I_d - UU^T$ the projector onto the orthogonal subspace, we have $(\Pi_U)^2 = \Pi_U$, $(\Pi_U^\perp)^2 = \Pi_U^\perp$, $\Pi_U^\perp U = 0$, $U^T \Pi_U^\perp = 0$.

The quotient geometry of $S^+(p, d)$ is thoroughly studied in [4]. Any tangent vector to $S^+(p, d)$ may be represented by the infinitesimal variation $(\delta U, \delta R) \in \mathcal{T}_{(U, R)} S^+(p, d)$ with $(\delta U, \delta R) \in \mathbb{R}^{d \times p} \times \mathbb{R}^{p \times p}$ of the following form

$$\delta U = (I_d - UU^T)\Gamma, \quad \Gamma \in \mathbb{R}^{d \times p}; \quad \delta R \in \mathbb{R}^{p \times p}, \quad \delta R^T = \delta R. \quad (6)$$

This may be interpreted as follows: An infinitesimal variation of U makes its columns move in the subspace of \mathbb{R}^d orthogonal to $\text{span}(U)$, as we have $(\delta U)^\top U = 0$, ensuring that $\delta(U^T U) = 0$, in accordance with the constraint $U^T U = I_p$. And an infinitesimal variation of R necessarily remains symmetric [4]. The corresponding tangent vector $\delta Y \in \mathcal{T}_Y S^+(p, d)$ at $Y = URU^T$ writes

$$\delta Y = (\delta U)RU^T + U(\delta R)U^T + UR(\delta U)^T. \quad (7)$$

2.2 Optimal approximations on the tangent space

Any tangent vector to the set of full $d \times d$ PSD matrices at X is encoded by a matrix $H \in \mathbb{R}^{d \times d}$, with H symmetric, see, e.g., [4]. Hence, given matrices $X(t)$ depending smoothly on the parameter $t \in \mathbb{R}$, one may write at each time $\dot{X}(t) = H$, and turn the problem (2), which specifies to (4) in the present case, into the generic problem of solving

$$\min_{\dot{Y}(t) \in \mathcal{T}_{Y(t)} S^+(p, d)} \|\dot{Y}(t) - H\|, \quad (8)$$

with H symmetric and arbitrary. Replacing H with $F(Y(t))$ at each time t , and completing it with an initial condition $Y(t_0)$ will yield an approximation $(Y(t))_{t \geq t_0}$ to the solution of the ODE $\dot{X} = F(X)$. We have the following:

Proposition 1 (from [16, 19]). *The orthogonal projection of a symmetric matrix H onto $\mathcal{T}_Y S^+(p, d)$ at $Y = URU^T$ is, in the retained form of tangent vectors (6),*

$$\delta Y = P_{U, R, s}(H) = \delta URU^T + U\delta RU^T + UR\delta U^T,$$

where the matrices are given by:

$$\delta U = (I - UU^T)HUR^{-1}, \quad (9)$$

$$\delta R = U^T H U. \quad (10)$$

The tangent vector then writes

$$P_{U,R,s}(H) = HU U^T + UU^T H - UU^T HUU^T.$$

This choice solves problem (8), that is, it minimizes over matrices of the form $\delta Y = \delta URU^T + U\delta RU^T + UR\delta U^T$ with constraints (6) the following cost

$$J_{H;U,R}(\delta U, \delta R) = \text{Tr}((H - \delta Y)^2) = \|H - \delta Y\|^2, \quad (11)$$

whose minimum is given by $\min J_{H;U,R}(\delta U, \delta R) = \text{Tr}((I - UU^T)H(I - UU^T)H)$, and where we used the trace form of Frobenius norm $\|A\|^2 = \text{Tr}(A^T A) = \text{Tr}(A^2)$ for symmetric matrices.

Given a differential equation $\frac{d}{dt}X(t) = F(X(t))$ over PSD matrices of full size $d \times d$, Equations (9)-(10) readily yield $\dot{U}(t)$ and $\dot{R}(t)$ to implement its low-rank approximation at all times, letting $H = F(Y(t))$ with $Y(t) = U(t)R(t)U(t)^T$.

3 Optimal low-rank plus diagonal approximation

We now turn to our main theoretical results which consist of extensions to the low-rank plus diagonal case.

3.1 Geometry of $S_{\text{diag}}^+(\mathfrak{p}, d)$ and $S_{\text{isot}}^+(\mathfrak{p}, d)$

Take $Y = URU^T + \psi$ with ψ diagonal and having all its diagonal elements strictly positive. Recalling (6), an element $Y \in S_{\text{diag}}^+(\mathfrak{p}, d)$, along with tangent vectors $\delta Y \in \mathcal{T}_Y S_{\text{diag}}^+(\mathfrak{p}, d)$ to this element, write

$$\left. \begin{aligned} Y &= URU^T + \psi \\ \delta Y &= \delta URU^T + U\delta RU^T + UR\delta U^T + \delta\psi \end{aligned} \right\} \quad (\text{FA}) \quad (12)$$

with $\delta U, \delta R$ as in (6) and $\delta\psi$ diagonal.

Remark 1. *In terms of geometric structure, $S^+(\mathfrak{p}, d)$ is a smooth embedded submanifold, albeit not straightforward to prove. The interested reader is referred to [31], which contains two proofs: Prop. 3.11, and another from [14] which is completed therein. Unfortunately, if we wonder whether this extends to $S_{\text{diag}}^+(\mathfrak{p}, d)$, we discover that $S_{\text{diag}}^+(\mathfrak{p}, d)$ is not a submanifold: The dimension of the tangent space is not constant, as can be seen by considering $Y = xx^T + \psi$ with $x = (1, 0, 0)^T$ on the one hand, leading to a dimension of 5 for all δY , and $x = (1, 1, 1)^T$ on the other, then leading to a dimension of 6. Thus, $S_{\text{diag}}^+(1, 3)$ bears no manifold structure, but this shall not prevent one from using (12) to represent the matrices and their derivatives, and refer to the set of all such δY as the tangent space at Y .*

Let us now turn to the PPCA decomposition, of the form $Y = URU^T + sI_d$. So far, it has been introduced as a subset of $S_{\text{diag}}^+(\mathfrak{p}, d)$ based on the constraint $\varphi = sI_d$, for the sake of simplicity of exposition. However, because some degrees of freedom are removed by this constraint, the eigenvalues of Y are necessarily lower bounded by $s > 0$, as $Y \succeq sI_d$, since URU^T is PSD. Thus, Y cannot have arbitrarily small eigenvalues in $\text{span}(U)$. This motivates the use of a slightly different—richer—parameterization. From now onwards, we define $S_{\text{isot}}^+(\mathfrak{p}, d)$ as matrices of the form

$$Y = U(R - sI)U^T + sI = URU^T + s(I - UU^T) = URU^T + s\Pi_U^\perp, \quad (13)$$

and impose that $R \in \mathbb{S}^+(\mathfrak{p}, \mathfrak{p})$, $U \in \mathbb{V}(\mathfrak{p}, \mathfrak{d})$ and $s > 0$. This allows for $\mathbb{S}_{\text{isot}}^+(\mathfrak{p}, \mathfrak{d})$ to be a bigger approximating set, while retaining that Y is positive definite.

This leads to the following geometry for $\mathbb{S}_{\text{isot}}^+(\mathfrak{p}, \mathfrak{d})$. An element $Y \in \mathbb{S}_{\text{isot}}^+(\mathfrak{p}, \mathfrak{d})$ with associated tangent vector $\delta Y \in \mathcal{TS}_{\text{isot}}^+(\mathfrak{p}, \mathfrak{d})$ in the PPCA form writes

$$\left. \begin{aligned} Y &= URU^T + s(I - UU^T) = U(R - sI)U^T + sI \\ \delta Y &= \delta U(R - sI)U^T + U(\delta R - \delta sI)U^T \\ &\quad + U(R - sI)\delta U^T + \delta sI \end{aligned} \right\} \text{(PPCA)} \quad (14)$$

with $\delta U, \delta R$ as in (6) and $\delta s \in \mathbb{R}$.

3.2 Optimal approximation in the PPCA form

Let us start with the simpler PPCA decomposition.

Proposition 2. *The orthogonal projection of a symmetric matrix $H \in \mathbb{R}^{d \times d}$ onto $\mathcal{T}_Y \in \mathbb{S}_{\text{isot}}^+(\mathfrak{p}, \mathfrak{d})$ is $\delta Y = \delta U(R - sI)U^T + U(\delta R - \delta sI)U^T + U(R - sI)\delta U^T + \delta sI$ where the matrices are given by:*

$$\begin{aligned} \delta s &= [\text{Tr}(H) - \text{Tr}(U^T H U)] / (d - p), \\ \delta U &= (I - UU^T) H U (R - sI)^{-1}, \\ \delta R &= U^T H U. \end{aligned} \quad (15)$$

The tangent vector then writes

$$P_{U,R,s}(H) = HUU^T + UU^T H - UU^T HUU^T + \delta s(I - UU^T),$$

This choice minimizes over matrices of the form (14) the following cost

$$C_{H;U,R,s}(\delta U, \delta R, \delta s) = \text{Tr}((H - \delta Y)^2). \quad (16)$$

Proof. The cost (16), that is, the squared Frobenius norm between H and the tangent vector, rewrites

$$\text{Tr}((\tilde{H} - [\delta U \tilde{R} U^T + U \delta R U^T + U \tilde{R} \delta U^T])^2),$$

with $\tilde{H} = H - \delta s(I - UU^T)$ and $\tilde{R} = R - sI$. In other terms

$$C_{H;U,R,s}(\delta U, \delta R, \delta s) = J_{\tilde{H};U,\tilde{R}}(\delta U, \delta R),$$

with J as in (11). The next step consists in writing

$$\min_{\delta U, \delta R, \delta s} C_{H;U,R,s} = \min_{\delta s} \left(\min_{\delta U, \delta R} C_{H;U,R,s} \right) \quad (17)$$

$$= \min_{\delta s} \left(\min_{\delta U, \delta R} J_{\tilde{H};U,\tilde{R}} \right) \quad (18)$$

$$= \min_{\delta s} \text{Tr}((I - UU^T) \tilde{H} (I - UU^T) \tilde{H}), \quad (19)$$

using Proposition 1, which also shows that in the present case $\min_{\delta U, \delta R} J_{\tilde{H};U,\tilde{R}}$ corresponds to

$$\delta U = (I - UU^T) \tilde{H} U \tilde{R}^{-1} = (I - UU^T) H U \tilde{R}^{-1} \quad (20)$$

$$\delta R = U^T \tilde{H} U = U^T H U. \quad (21)$$

To conclude, we still have to solve (19). Recalling $\tilde{H} = H - \delta s(I - UU^T)$ we differentiate $\text{Tr}((I - UU^T) \tilde{H} (I - UU^T) \tilde{H})$ with respect to δs

$$\begin{aligned} J'(\delta s) &= -2 \text{Tr}((I - UU^T) \tilde{H}) \\ &= -2 \text{Tr}((I - UU^T)(H - \delta s(I - UU^T))) \end{aligned}$$

and we find $\delta s = [\text{Tr}(H) - \text{Tr}(U^T H U)] / (d - p)$. \square

3.3 Optimal approximation in the FA form

The problem of projecting onto $S_{\text{diag}}^+(\mathfrak{p}, \mathfrak{d})$ is more difficult than the latter, as in the proof of Proposition 2 we extensively used the isotropy of the diagonal term sI_d , namely $U(sI_d)U^T = sUU^T$. However, it turns out we may find a closed-form expression also in the FA decomposition. This allows for capturing different individual variances in the subspace orthogonal to $\text{span}(U)$.

The first step is to transform the optimization problem over the matrix ψ into a least-squares problem for a vector consisting of its diagonal. The cost we want to minimize is

$$\begin{aligned} C_{H;U,R,\psi}(\delta U, \delta R, \delta\psi) &:= \text{Tr}((H - [\delta URU^T + U\delta RU^T + UR\delta U^T + \delta\psi])^2) \\ &= \text{Tr}((H - \delta\psi - [\delta URU^T + U\delta RU^T + UR\delta U^T])^2) \\ &= \text{Tr}((\tilde{H} - [\delta URU^T + U\delta RU^T + UR\delta U^T])^2) \\ &= J_{\tilde{H},U,R}(\delta U, \delta R), \end{aligned}$$

where now \tilde{H} denotes $\tilde{H} = H - \delta\psi$. We write

$$\min_{\delta U, \delta R, \delta\psi} C(\delta U, \delta R, \delta\psi) = \min_{\delta\psi} \left(\min_{\delta U, \delta R} J_{\tilde{H};U,\tilde{R}}(\delta U, \delta R) \right),$$

which is equal to $\min_{\delta\psi} J_{\tilde{H};U,R}(\delta U, \delta R)$ with

$$\delta U = (I - UU^T)\tilde{H}UR^{-1}, \quad (22)$$

$$\delta R = U^T\tilde{H}U, \quad (23)$$

using (9), (10). Using Proposition 1, we see the cost function to minimize becomes

$$\begin{aligned} \tilde{C}(\delta\psi) &:= \text{Tr}((I - UU^T)\tilde{H}(I - UU^T)\tilde{H}) \\ &= \text{Tr}(\tilde{H}^2) - 2\text{Tr}(UU^T\tilde{H}^2) + \text{Tr}(UU^T\tilde{H}UU^T\tilde{H}). \end{aligned} \quad (24)$$

(24) may be re-written as a vector least squares problem for the vector $\overline{\delta\psi} := (\delta\psi_{11}, \delta\psi_{22}, \dots, \delta\psi_{dd})^T = \text{diag}(\delta\psi) \in \mathbb{R}^d$. This leads to the following optimal approximation.

Proposition 3. *The orthogonal projection of a symmetric matrix $H \in \mathbb{R}^{d \times d}$ onto $\mathcal{T}_Y \in S_{\text{diag}}^+(\mathfrak{p}, \mathfrak{d})$ is $\delta Y = \delta URU^T + U\delta RU^T + UR\delta U^T + \delta\psi$ where the matrices are given by*

$$\begin{aligned} \text{diag}(\delta\psi) &= \overline{\delta\psi}, \\ \delta U &= (I - UU^T)(H - \delta\psi)UR^{-1}, \\ \delta R &= U^T(H - \delta\psi)U, \end{aligned} \quad (25)$$

and where the vector $\overline{\delta\psi} \in \mathbb{R}^d$ that encodes the diagonal matrix $\delta\psi \in \mathbb{R}^{d \times d}$ is defined by

$$\overline{\delta\psi} = ((I - UU^T) \circ^2)^+ \text{diag}((I - UU^T)H(I - UU^T)), \quad (26)$$

letting $+$ denote the Moore-Penrose inverse and where \circ denotes the element-wise (Hadamard) matrix product.

Proof. Recall the notation $\Pi_U^\perp = I_d - UU^T$. We may rewrite the cost (24) in terms of the vector $\overline{\delta\psi}$, i.e., $\tilde{C}(\delta\psi) = \bar{C}(\overline{\delta\psi})$ and we find:

$$\begin{aligned}\bar{C}(\overline{\delta\psi}) &= \text{Tr}((\Pi_U^\perp H - \Pi_U^\perp \text{diag}(\overline{\delta\psi}))^2) \\ &= c - 2\text{Tr}(\Pi_U^\perp H \Pi_U^\perp \text{diag}(\overline{\delta\psi})) + \text{Tr}(\text{diag}(\overline{\delta\psi}) \Pi_U^\perp \text{diag}(\overline{\delta\psi}) \Pi_U^\perp),\end{aligned}$$

where c is a constant with respect to $\overline{\delta\psi}$. This may be re-written as a vector least-squares problem for the vector $\overline{\delta\psi}$ using the relation $\text{Tr}(\text{diag}(x)A\text{diag}(y)B^T) = x^T(A \circ B)y$ where \circ is the element-wise matrix product. If we note $\mathbf{1}_d$ the vector of ones, the cost becomes:

$$\tilde{C}(\delta\psi) = c - 2\mathbf{1}_d^T (\Pi_U^\perp H \Pi_U^\perp \circ I) \overline{\delta\psi} + \overline{\delta\psi}^T (\Pi_U^\perp \circ \Pi_U^\perp) \overline{\delta\psi}.$$

Taking the derivative with respect to $\overline{\delta\psi}$, we obtain:

$$\nabla \bar{C}(\overline{\delta\psi}) = -2\text{diag}(\Pi_U^\perp H \Pi_U^\perp) + 2(\Pi_U^\perp \circ \Pi_U^\perp) \overline{\delta\psi}$$

Zeroing the gradient yields (26) indeed. □

4 Implementation

We now discuss the interest of the results for the high-dimensional setting. Our goal, for tractability, is to elicit operations and storage costs being linear in the dimension d . We prove that both methods lead to projections that require linear computation cost in the dimension d , and discuss the numerical cost in detail. A key result to our analysis is that any matrix Y of $S_{\text{diag}}^+(\mathbf{p}, d)$ or $S_{\text{isot}}^+(\mathbf{p}, d)$ is easily inverted via the Woodbury lemma:

$$(URU^T + \psi)^{-1} = \psi^{-1} - \psi^{-1}U(R^{-1} + U^T\psi^{-1}U)^{-1}U^T\psi^{-1}, \quad (27)$$

a fact we will extensively use. Note that it also proves matrices of $S_{\text{diag}}^+(\mathbf{p}, d)$ and $S_{\text{isot}}^+(\mathbf{p}, d)$, contrary to low-rank matrices of $S^+(\mathbf{p}, d)$, can be inverted, and this with a numerical cost of inversion is $O(p^3d)$, hence linear in the dimension d .

4.1 Numerically efficient formulation for the FA form

Equation (26) necessitates (pseudo)-inversion of a $d \times d$ matrix $(I - UU^T)^{\circ 2}$ which may hinder its use in large dimension. However, it is amenable to a linear computation cost.

Proposition 4. *The orthogonal projection (26) can be performed with linear computation cost in d . Indeed, the required vector $\overline{\delta\psi} \in \mathbb{R}^d$ is given by (30) below, and can be advantageously computed via (36).*

To prove the result we start with a lemma, whose proof is technical and hence postponed to the Appendix.

Lemma 1. *The cost $\tilde{C}(\delta\psi)$ may be re-written as a function of the vector $\overline{\delta\psi} \in \mathbb{R}^d$ as follows:*

$$\begin{aligned}\tilde{C}(\delta\psi) &= \bar{C}(\overline{\delta\psi}) := \alpha + (\bar{h} - \overline{\delta\psi})^T (\bar{h} - \overline{\delta\psi}) \\ &\quad + 4\bar{h}_U^T \overline{\delta\psi} - 2\overline{\delta\psi}^T \bar{D} \overline{\delta\psi} - 2\Lambda^T \overline{\delta\psi} + \overline{\delta\psi}^T \Upsilon \Upsilon^T \overline{\delta\psi},\end{aligned} \quad (28)$$

where α is a constant, $\bar{h}, \bar{h}_U, \Lambda$ are all vectors of \mathbb{R}^d , and were \bar{D} is a diagonal $d \times d$ matrix, and Υ is a matrix, all being given below. Hence the gradient writes

$$\frac{1}{2} \nabla \bar{C}(\bar{\psi}) = (I - 2\bar{D} + \Upsilon \Upsilon^T) \bar{\psi} - \Lambda + 2\bar{h}_U - \bar{h}, \quad (29)$$

and the optimizer is given by

$$\bar{\psi} = (I - 2\bar{D} + \Upsilon \Upsilon^T)^+ (\bar{h} - 2\bar{h}_U + \Lambda), \quad (30)$$

The parameters in the equations above are:

$$\bar{h} = \text{diag}(H), \quad \bar{h}_U := \text{diag}(UU^T H) \quad (31)$$

$$\Lambda = (\Lambda_1, \dots, \Lambda_d)^T \in \mathbb{R}^d \quad (32)$$

$$\Lambda_k = \sum_{1 \leq i, j \leq r} (U^T H U)_{ij} U_{ki} U_{kj}, \quad 1 \leq k \leq d \quad (33)$$

$$\Gamma_{ij} = ((\Gamma_{ij})_1, \dots, (\Gamma_{ij})_d)^T \in \mathbb{R}^d \quad (34)$$

$$(\Gamma_{ij})_k = U_{ki} U_{kj}, \quad 1 \leq k \leq d, \quad (35)$$

with Υ a matrix of size $d \times \frac{p(p+1)}{2}$ whose first $\frac{p(p-1)}{2}$ columns are given by the vectors $\sqrt{2}\Gamma_{ij}$, $i < j$, and whose last r columns are Γ_{ii} , and \bar{D} is the diagonal matrix defined by $\bar{D}_{ii} = \sum_{j=1}^r U_{ij}^2$. It may be checked (30) coincides with (26).

The problem with (30) is that it requires inverting a $d \times d$ matrix, which is not affordable computationally. When $I - 2\bar{D} + \Upsilon \Upsilon^T$ is invertible, we can use Woodbury matrix identity instead to express the solution (30) as

$$\bar{\psi} = (\varphi - \varphi \Upsilon (I_{\frac{p(p+1)}{2}} + \Upsilon^T \varphi \Upsilon)^{-1} \Upsilon^T \varphi) (\bar{h} - 2\bar{h}_U + \Lambda), \quad (36)$$

where $\varphi := (I_d - 2\bar{D})^{-1}$ is the inverse of a diagonal matrix, which may be efficiently computed using only vectors of \mathbb{R}^d , and where the $\frac{p(p+1)}{2} \times \frac{p(p+1)}{2}$ matrix $I_{\frac{p(p+1)}{2}} + \Upsilon^T \varphi \Upsilon$ needs be inverted.

4.2 Computation cost

As a preliminary remark, we see that whatever the chosen low-rank decomposition, computing the approximation requires storing and projecting the $d \times d$ symmetric matrix H . To perform this in a high dimension we need to assume some sort of sparsity, or low-rank structure for H of the kind $H = GG^T$ with $G \in \mathbb{R}^{d \times r}$ where $r \ll d$, or $H \in S_{\psi}^+(r, d)$. This problem is already present in the literature on low-rank approximations and has proved not to be a limitation in applications, see [16, 19] and articles that followed. For our complexity analysis, we will henceforth suppose the computational cost of computing $U^T H U$ is $O(dp^2)$. Note that to project the matrix in the PPCA or FA form, all equations must be coded ‘‘vectorially’’, that is, diagonal matrices are encoded by vectors, and we carefully choose the order of operations to never multiply full-size matrices, e.g., in (36).

Propositions 2, 3 and 4 give the formulas that one needs to evolve the factors $U(t), R(t), s(t)$ or $U(t), R(t), \psi(t)$, where their derivatives are given by (15) and (25), (36) respectively. In the PPCA form, to compute matrices (15) we see we need to perform $O(dp^2)$ operations. In the FA form, implementation of (25) retains linear complexity in d . However in the efficient implementation (36) of (26), we need to invert the $\frac{p(p+1)}{2} \times \frac{p(p+1)}{2}$ matrix $I_{\frac{p(p+1)}{2}} + \Upsilon^T \psi \Upsilon$. This yields a computational

cost of order $O(dp^2) + O(p^6)$. We thus see that FA decomposition—albeit richer—comes at a price: the number of “factors” p that may be used is more limited than in the PPCA decomposition.

To illustrate how the methods compare in terms of numerical cost, we propose to apply the three methods for the approximation of a large-scale matrix $H = GG^T$ where G is a matrix of size $d \times r$ generated randomly, where $r > p$ is low, but sufficiently large for the p -rank approximation not to be exact. Using a product form for H allows for low-memory-cost operations, in particular, we have $UHU^T = (UG)(UG)^T$ yielding a linear cost in d . The execution times of the different algorithms are shown in Table 1 in dimension $d = 10^6$ where the computations are performed on a standard laptop with Octave. All algorithms scale well to very high-dimensional problems. However, FA has a higher computation time since it needs to invert a $\frac{p(p+1)}{2} \times \frac{p(p+1)}{2}$ matrix.

Table 1: Matrix projection with $d = 10^6$, $p = 10$, and $r = 100$

Set	Projection method	Execution time (s)	Memory cost
$S^+(p, d)$	“Low-rank”	3.5 s	dp
$S_{\text{isot}}^+(p, d)$	“PPCA”	7 s	dp
$S_{\text{diag}}^+(p, d)$	“FA”	35 s	$dp + p^4/4$

5 Application to the Riccati equation

In this section, we consider the following classical continuous-time Riccati equation in high dimension d :

$$\frac{d}{dt}P = AP + PA^T + Q - PC^T N^{-1}CP, \quad (37)$$

with P , A and Q are $d \times d$ high dimensional matrices, N is a $k \times k$ matrix and C is of size $k \times d$. P , Q and N are PSD matrices.

To store this covariance matrix P in high dimension a first solution is to project it onto the low-rank manifold $S^+(p, d)$. To do so we replace P with $Y = URU^T$ in the right-hand side (37), and we project the obtained matrix. Prop. 1 yields

$$\begin{aligned} \dot{U} &= (I - UU^T)(AU + QUR^{-1}) \\ \dot{R} &= U^T AUR + RU^T A^T U + U^T QU - RU^T C^T N^{-1}CUR, \end{aligned} \quad (38)$$

termed low-rank Riccati equation in [7]. This may be used as a low-rank proxy to the Kalman filter. The obtained filter then performs inference in the low-dimensional subspace spanned by the dominant eigenvectors [20].

To approximate (37) in $S_{\text{isot}}^+(p, d)$ instead, we replace P with $Y = URU^T + s(I - UU^T)$ in the right-hand side of (37) and we project the obtained matrix using (15).

Proposition 5 (PPCA-Riccati). *The lift on $S_{\text{isot}}^+(p, d)$ of the orthogonal projection of the vector field defined by (37) writes*

$$\dot{s} = \frac{1}{d-p} \text{Tr}((I - UU^T)(2sA + Q - s^2 C^T N^{-1}C)) \quad (39)$$

$$\dot{U} = (I - UU^T)(AUR + QU + sA^T U - sC^T N^{-1}CUR)(R - sI)^{-1} \quad (40)$$

$$\dot{R} = U^T AUR + RU^T A^T U + U^T QU - RU^T C^T N^{-1}CUR. \quad (41)$$

We have the following interpretation: (41) is the original Riccati equation (37) where matrices are projected onto the subspace encoded by U , (40) resembles the Oja flow [24] that tracks the dominant subspace of a symmetric matrix, and (39) provides an adaptation of parameter s to reflect the inflation of the covariance under the vector field in the subspace orthogonal to $\text{span}(U)$. Note that, (39) offers guarantees that $s(t)$ remains positive at all times, since as long as $(I - UU^T)Q \neq 0$, we have $s = 0 \Rightarrow \dot{s} > 0$. Thus the lifted equations (39)-(41) preserve our PPCA form at all times.

In terms of implementation, we will need a numerical integration method for the obtained differential equations. With h the time step and $\dot{s}(t) = \delta s, \dot{U}(t) = \delta U, \dot{R}(t) = \delta R$, one needs to define $s(t+h), U(t+h), R(t+h)$ so as $Y(t+h) \in \text{S}_{\text{isot}}^+(p, d)$. This can be performed using a retraction, see [1]. We let

$$s(t+h) \leftarrow s(t) + h\delta s \quad (42)$$

$$U(t+h) \leftarrow \text{retraction_qr}(U(t), h\delta U) \quad (43)$$

$$R(t+h) \leftarrow R^{1/2} \exp(hR^{-1/2}\delta R R^{-1/2})R^{1/2}, \quad (44)$$

where (43) is based on the retraction “retraction_qr” of the Stiefel manifold from the Manopt toolbox [6], and (44) is a natural retraction on PD matrices [4], and where “exp” denotes the usual matrix exponential, whose computation is tractable in the small dimension p . $\text{S}_{\text{diag}}^+(p, d)$ is similarly treated. Note that alternative schemes have been proposed on $\text{S}^+(p, d)$, see [16, 19, 23]. [22] even proposes a numerical integration method that avoids inverting the factor R , see [8] for recent developments.

Regarding the FA approach, we can approximate (37) in $\text{S}_{\text{diag}}^+(p, d)$ as follows. We write $H := \dot{P} = AP + PA^T + Q - PSP$ with $S = C^T N^{-1} C$ and $P = URU^T + \psi$. After some calculations this yields

Proposition 6 (FA-Riccati). *The lift on $\text{S}_{\text{diag}}^+(p, d)$ of the orthogonal projection of the vector field defined by (37) writes*

$$\begin{aligned} \text{diag}(\dot{\psi}) &= (I - 2\bar{D} + \Upsilon\Upsilon^T)^+ \text{diag}((I - UU^T)M(I - UU^T)) \\ \dot{U} &= (I - UU^T)(M - \delta\psi)UR^{-1} + (I - UU^T)(AU - \psi SU) \\ \dot{R} &= U^T(M - \delta\psi)U - U^T\psi SUR - RU^T\psi SU - RU^T SUR + U^T AUR + RU^T A^T U, \end{aligned}$$

where $M = A\psi + \psi A^T + Q - \psi S\psi$, $S = C^T N^{-1} C$ and \bar{D}, Υ as in (30).

The pseudo-inversion may be avoided when $p \ll d$ using the Woodbury formula as in equation (36), leading to an update having a memory cost linear in d if the order of operations is taken carefully.

We now propose two applications for the equations just obtained, one in the field of computational statistics where one seeks to approximate a fixed posterior distribution, for which we provide a tractable approximation in high dimension through a Riccati flow, and one in the field of Kalman filtering, with applications to robotics. The former is essentially an illustration of the applicability of the method, while the latter aims at comparing the various approaches through numerical simulations.

5.1 Application to computational statistics

We seek to approximate a distribution in high dimension which is the stationary solution of a Langevin equation. In this context, Markov-Chain Monte Carlo (MCMC) has been extensively

studied in statistical physics and machine learning [2]. Recently, variational inference [18] has been used in this framework as an alternative form to MCMC to approximate the stationary distribution with a Gaussian distribution. The final approximation is defined as the asymptotic limit of a flow over the set of Gaussian distributions, parameterized by their mean and covariance matrix. The flow for the covariance matrix has the form of a Riccati equation. To illustrate the applicability of the methods developed in the present paper, we now show that the PPCA projection can be directly applied to these flows to provide a memory-efficient and computationally tractable algorithm in high dimension.

5.1.1 Wasserstein gradient flow for Gaussian variational inference

In Bayesian statistics, one is often faced with the problem of approximating a target posterior distribution π that is known up to a normalizing constant, that is, we have access to $\nabla V(x)$ where $\pi(x) \propto \exp(-V(x)/\epsilon)$. To compute a Gaussian approximation of π we may proceed as follows. The following stochastic differential equation (Langevin dynamics):

$$dx_t = -\nabla V(x_t)dt + \sqrt{2\epsilon}dB_t, \quad (45)$$

is such that, under suitable assumptions, the marginal distribution of x_t satisfies the Fokker-Planck (partial differential) equation:

$$\frac{\partial p_t}{\partial t} = \operatorname{div}(Vp_t) + \epsilon\Delta p_t, \quad (46)$$

which has stationary distribution π . Thus, if one can approximate the solution to the latter equation by a Gaussian density, one may hope the latter tends asymptotically to a Gaussian approximation of π . In the recent work [18], the solution of this PDE has been approximated with a Gaussian $p_t \approx q_t = \mathcal{N}(\mu_t, P_t)$ using variational inference. The following Gaussian flow:

$$\begin{aligned} \dot{\mu}_t &= -\mathbb{E}_{X \sim \mathcal{N}(\mu_t, P_t)}[\nabla V(X)], \\ \dot{P}_t &= A_t P_t + P_t A_t^T + Q, \\ Q &:= 2\epsilon I; \quad A_t := -\mathbb{E}_{X \sim \mathcal{N}(\mu_t, P_t)}[\nabla^2 V(X)], \end{aligned} \quad (47)$$

is shown to converge with exponential rate to the Gaussian distribution being the closest to the target distribution π (in the sense of KL divergence), if V is strongly convex, see [18, Appendix A]. For large-scale problems, the differential equation on P_t in equation (47) is problematic because it requires storing at each time a $d \times d$ matrix. Moreover, the computation cost is (roughly) cubic in the dimension due to matrix products, even if we overlook the difficulty of computing the required expectations. This motivates a low-rank plus diagonal approximation to this Riccati-like equation. We readily see the Wasserstein gradient flow (47) lends itself to the present low-rank approximation framework, as it is of the form (37), provided that the expectations which define A_t may be computed, and that they can be written in a factorized form compatible with the high dimension.

5.1.2 Particular case of a Gaussian target

To fix ideas, let us see what the equations boil down to in the case where the target is Gaussian, that is, $\pi \sim \mathcal{N}(m, \epsilon M)$, to be consistent with $\pi \propto \exp(-V/\epsilon)$ so that $V(x) = \frac{1}{2}(x - m)^T M^{-1}(x - m)$.

In this case, the expectations can be computed analytically and the variational Gaussian flow (47) writes:

$$\begin{aligned}\dot{\mu}_t &= -M^{-1}(\mu_t - m) \\ \dot{P}_t &= 2\varepsilon I - M^{-1}P_t - P_tM^{-1}.\end{aligned}$$

The equations have as stationary point $(\mu, P) = (m, \varepsilon M)$ which are the parameters of π , as expected.

Remark 2. We see the “pure” low-rank projection, defined by equation (9), of the covariance P_t writes:

$$\dot{U} = -(I - UU^T)M^{-1}U,$$

and we recover minus the Oja flow [24], that is, the flow in a form that tracks the eigenvectors corresponding to the smallest eigenvalues of M^{-1} , thus, the dominant eigenspace of M .

5.1.3 PPCA approximation in the large-scale case

We now come back to a general target posterior $\pi \propto \exp(-V/\varepsilon)$ and seek to approximate the Gaussian flow (47) in a tractable form with high d . The first step is to approximate the expectation under the Gaussian distribution appearing in the Gaussian flow (47) with K Monte-Carlo samples. To this aim, we may let appear an outer product of the form:

$$\begin{aligned}\mathbb{E}_{\mathcal{N}(\mu_t, P_t)}[\nabla^2 V(X)]P_t &= \mathbb{E}_{\mathcal{N}(\mu_t, P_t)}[\nabla V(X)(X - \mu_t)^T] \\ &\approx \frac{1}{K} \sum_{k=1}^K \nabla V(x_t^k)(x_t^k - \mu_t)^T = D_t B_t^T \quad \text{where } x_t^k \sim q_t, \\ D_t &:= \frac{1}{\sqrt{K}} (\nabla V(x_t^1), \dots, \nabla V(x_t^K)) \quad B_t := \frac{1}{\sqrt{K}} (x_t^1 - \mu_t, \dots, x_t^K - \mu_t),\end{aligned}$$

where the first equality comes from integration by parts (Stein Lemma [21]). We now only need to store the matrices B_t and D_t of size $d \times K$. In practice, the sampling is done after the discretization of these ODEs at each integration step.

As D_t and B_t depend on the samples x_t^k generated from the distribution $q_t \sim \mathcal{N}(\mu_t, P_t)$, where in our PPCA approximation we let $P_t = URU^T + s(I - UU^T)$, we need to be able to sample from a Gaussian distribution having such a factorized covariance matrix. A simple way to do so is as follows. We first note that if we take $Z_1 \sim \mathcal{N}(0, URU^T)$ and $Z_2 \sim \mathcal{N}(0, s(I - UU^T))$ independent, then $Z_1 + Z_2$ has the desired covariance. To sample Z_1 , one may sample a small-dimensional variable $Z_3 \sim \mathcal{N}(0, R)$ and let $Z_1 = UZ_3$. To sample Z_2 , we may sample $Z_4 \sim \mathcal{N}(0, sI)$ and let $Z_2 = (I - UU^T)Z_4$.

The ODE we seek to approximate now takes the form:

$$\dot{P}_t = 2\varepsilon I - (D_t B_t^T + B_t D_t^T). \tag{48}$$

We project this equation onto the PPCA subset. This is achieved thanks to Proposition 5, and it yields

$$\begin{aligned}\dot{s} &= \frac{1}{d-p} \text{Tr}(2(I - UU^T)(\varepsilon I - D_t B_t^T)), \\ \dot{U} &= (I - UU^T)(2\varepsilon I - D_t B_t^T - B_t D_t^T)U(R - sI)^{-1}, \\ \dot{R} &= U^T(2\varepsilon I - D_t B_t^T - B_t D_t^T)U.\end{aligned}$$

These operations can be computed in a memory-efficient way using the relation $\text{Tr}D_t B_t^T = \sum_{i=1}^d D_t[i, :] B_t[:, i]$ and $U^T D_t B_t^T U = (U^T D_t)(U^T B_t)^T$ to only manipulate vectors or matrices of size $p \times K$.

We have briefly shown that using the results of the present paper, the recent Wasserstein gradient flow of [18] for variational inference is amenable to a PPCA approximation being compatible with the high dimension. We now turn to another application and compare numerically the various low-rank approaches.

5.2 Application to Kalman filtering

In this section, we consider the more standard problem of the Kalman filter in continuous time known as the Kalman-Bucy filter. In Kalman filtering, we seek to estimate hidden physical quantities that evolve over time and which are partially observed through a linear model. The covariance of the state P satisfies the Riccati equation (37).

To assess the various low-rank approximations, we consider an example inspired by robotics. We consider a swarm of $d/2$ agents in the 2D plane, each equipped with motion sensors, and governed by the following dynamics

$$\forall 1 \leq i \leq d/2 \quad \frac{d}{dt} X_i = u_i(t) + w_i(t), \quad (49)$$

where $X_i \in \mathbb{R}^2$ denotes the position of agent i , $w_i(t)$ is a white noise, and $u_i(t)$ a control input. Measurements consist of relative position between some agents being neighbors in a visibility graph, corrupted by noise $v_{ij}(t)$, i.e.,

$$y_{ij}(t) = X_j(t) - X_i(t) + v_{ij}(t), \quad (i, j) \in \Gamma. \quad (50)$$

Moreover, there is a "queen", say, agent 1, having relatively more computational capacity onboard—albeit limited—and which is equipped with a GPS, that is, we also measure $y_1(t) = X_1(t) + v_1(t)$. The queen receives the dynamical motions $u_i(t)$ as well as the measurements $y_{ij}(t)$, and estimates the state of the whole swarm onboard, through a low-rank filter compatible with its modest computational capabilities.

We integrate the Riccati equation with an Euler scheme (with step $0.01s$) during $10s$ and consider an initial covariance matrix factorized as purely low-rank $P_0 = URU^T$ with $R = 2I_p$ and U a random Stiefel matrix, to allow for a common starting point. We set noise covariance to be $N = 2I$ and process noise covariance $Q = D$ where D is a diagonal matrix whose diagonal consists of positive values dispersed around 1, reflecting discrepancies in the accuracy of motion sensors of the agents. In our experiments, each agent sees one other agent (that might be the queen), randomly picked at the beginning. We tested changing the visibility graph over time but that does not change the nature of the following results.

We compare the low-rank approximation [16, 19] with the two variants of our low-rank + diagonal approximation in dimension $d = 200$, that is, 100 agents, and a latent dimension $p = 8$ or $p = 50$ in Figure 1. We clearly observe the proposed approximations outperform the former in terms of distance to the true covariance matrix. Moreover, the order of the curves is as expected: Projections onto an increasing sequence of submanifolds yield in turn increasing accuracy for the matrix differential equation approximation.

To assess the effect of approximating the Riccati equation on the Kalman filter's state estimates, we have compared the full-rank KF's optimal estimates with those obtained by its computationally cheaper variants, for a randomly distributed initial error, in a noise-free setting, that is, when using

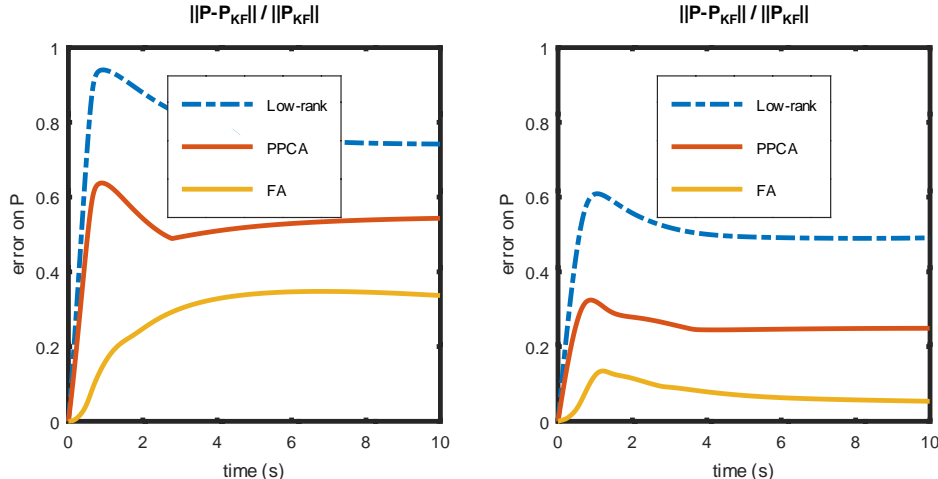


Figure 1: $d = 200$ with $p = 8$ (left) and $p = 50$ (right). Normalized distance between the covariance matrix computed from the true full-rank Riccati equation and the ones computed from the low-rank, low-rank + diagonal (FA) and low-rank + isotropic diagonal (PPCA) for the swarm example.




Figure 2: $d = 200$ with $p = 8$ (left) and $p = 50$ (right). Norm of the error over time between the filters' state estimates and the full KF's optimal estimate, for a randomly picked initial error in a noise-free setting.

the Kalman filter as an observer, to get more legible curves. The results are given in Figure 2. We see the deviation to the optimal estimates is much more contained when adding a diagonal matrix.

It is striking to see the PPCA method with $p = 8$ competes with the low-rank method of [19] (arxiv preprint version) with $p = 50$ in terms of covariance matrix approximation, at the expense of only 1 more scalar parameter. We also observe PPCA with $p = 8$ outperforms low-rank with $p = 50$ in terms of state estimate accuracy. More generally, the experiments show that adding a diagonal term provides an efficient alternative to increasing the rank.

Two remarks are in order. First, we note that, although the curves indicate a clear improvement of FA over PPCA, this is in fact largely due to the choice of anisotropic process noise, reflecting a discrepancy in the motion sensors' accuracy. When making the problem more "isotropic", the differences between the two diminish. Then, we also observed that the initial condition $P(0)$ plays

an important role. Although we took $P(0) \in S^+(p, d)$ to allow for a common starting point, we noted that departing from $P(0) \in S_{\text{isot}}^+(p, d)$ naturally reduces the gap between FA and PPCA.

5.2.1 Source code

The code is available on Github for Octave or Matlab at <https://github.com/marc-h-lambert/Riccati-PPCA>. It provides projections onto $\mathcal{TS}^+(p, d)$, $\mathcal{TS}_{\text{isot}}^+(p, d)$ and $\mathcal{TS}_{\text{diag}}^+(p, d)$ of matrices of the form $H = GG^T$, or diagonal, or a linear combination, and scripts to redo the numerical experiments.

5.2.2 Interest of invertible approximations in Kalman filtering

We conclude this section with a discussion about the relevance of low-rank approximations in Kalman filtering. To illustrate the shortcomings of purely low-rank based methods, e.g., [19, 30, 5, 16, 20] for Kalman filtering, we consider a tutorial case, see Example 6.2.10 of [25]. Consider noisy observations of a Brownian motion in a large dimension

$$dX = dw, \quad dY = dX + dv, \quad P(t_0) = \bar{P},$$

with w, v Wiener process noises with covariance matrices $Q = \lambda I$, $N = \nu I$. Assume one wants to filter the noise out.

Low-rank KF The low-rank Riccati equation of [4, 19, 20] writes

$$\dot{U} = 0, \quad \dot{R} = U^T Q U - R U^T N^{-1} U R$$

as we have $\dot{U} = (I - U U^T) Q U R^{-1} = \lambda (I - U U^T) U R^{-1}$. As a result, R stabilizes to the steady-state solution $\sqrt{\lambda \nu} I_p$. Let us see how the corresponding steady-state Kalman filter updates the state. The static Kalman gain writes $K = P_\infty C^T N^{-1} = \sqrt{\lambda/\nu} U_0 U_0^T$ where U_0 is the initial value of U . The steady-state KF equations are

$$d\hat{X} = K(dY - \hat{X} dt) = \sqrt{\lambda/\nu} U U^T (dY - \hat{X} dt).$$

We see that P is a degenerate low-rank matrix and the innovation vector $dY - \hat{X} dt$ is always projected onto the same subspace by the operator $U U^T$.

Low-rank plus diagonal KF in PPCA form By contrast, (39)-(41) write:

$$\dot{U} = 0, \quad \dot{R} = \lambda I_p - \frac{1}{\nu} R^2, \quad \dot{s} = \lambda - \frac{s^2}{\nu}.$$

The steady-state solution gives $R = \sqrt{\lambda \nu} I_p$ and $s = \sqrt{\lambda \nu}$. The associated KF equations write now:

$$P = U R U^T + s(I - U U^T) = \sqrt{\lambda \nu} I$$

$$d\hat{X} = K(dY - \hat{X} dt) = \sqrt{\lambda/\nu} (dY - \hat{X} dt).$$

In this particular example, we see we recover the equations of the full-rank steady-state Kalman filter.

Proposition 7. *Using the low-rank KF, the estimation error incurred grows unbounded $E(\|X - \hat{X}\|^2) \rightarrow +\infty$. By contrast, using our low-rank plus diagonal approximation, $E(\|X - \hat{X}\|^2) \rightarrow \sqrt{\lambda/\nu d}$. Hence, the error is here optimally contained.*

Proof. For the low-rank KF we have $d[(I - UU^T)(X - \hat{X})] = (I - UU^T)dw$. Thus $\frac{d}{dt}E(\|X - \hat{X}\|^2) \geq (d - p)\lambda$. By contrast, for the proposed KF, the estimation error is an Ornstein-Uhlenbeck process whose norm stabilizes at $\sqrt{\lambda/\nu d}$. \square

Contrary to the robotics example, the latter problem obviously lends itself to the proposed approximation, owing to C, Q, N being diagonal and isotropic, leading to a diagonal and isotropic (stationary) covariance that is fully recovered on $S_{\text{isot}}^+(p, d)$. However, it clearly illustrates the interest of maintaining an invertible covariance matrix in Kalman filtering.

6 Conclusion

We have derived the orthogonal projection of any tangent vector to the set of PSD matrices onto the tangent space to low-rank plus diagonal PSD matrices. This allows for an attractive alternative to previous low-rank approximations of differential equations defined on the set of PSD matrices, in that it retains their computational efficiency while allowing for more flexible, full-rank, approximations, leading to better results for large-scale filtering. Our conclusion is as follows:

1. Existing (pure) low-rank PSD approximations have the shortcoming of having many null eigenvalues, which comes with consequences. If they model a Gaussian covariance matrix, we cannot compute the precision matrix, and there is no density. In Kalman filtering, this results in overconfidence, potentially leading to large error drift in turn.
2. The proposed PPCA approximation is about as efficient in terms of memory and computational cost while having none of those drawbacks. We recommend its systematic use.
3. The proposed FA approximation additionally closely captures the diagonal elements. However, it comes at a greater price, albeit retaining linearity in d . Its use shall be reserved for problems known to have a strongly non-isotropic diagonal, which one wants to capture well.

As a perspective, we see that while PPCA is computationally efficient, it doesn't reach the error accuracies achieved by the FA format. It could prove useful to seek a practical criterion for switching between these two formats. Besides, a criterion for choosing the rank p (or adapting it dynamically over time) may also prove useful. A first step to address those perspectives is to observe that the approximation error $\|H - P(H)\|^2$ can be computed at all times with a cost being linear in d too, allowing for approximation quality assessment, as shown in Appendix B. An other interesting perspective would be to explore how one could bring the proposed method to bear on the dynamic factor analysis problem, see e.g., [12].

Acknowledgments

This work was funded by the French Defence procurement agency (DGA) and by the French government under the management of Agence Nationale de la Recherche as part of the "Investissements d'avenir" program, reference ANR-19-P3IA-0001(PRAIRIE 3IA Institute). We also acknowledge support from the European Research Council (grant SEQUOIA 724063).

We thank Pierre-Antoine Absil and Guillaume Ollier for enlightening discussions about the submanifold structure of $S^+(p, d)$. We also thank the anonymous reviewers for their suggestions.

References

- [1] Pierre-Antoine Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization Algorithms on Matrix Manifolds*, volume 78. Princeton University Press, 2008.
- [2] Christophe Andrieu, Nando de Freitas, Arnaud Doucet, and Michael I. Jordan. An introduction to mcmc for machine learning. *Machine Learning*, 50(1):5–43, 2003.
- [3] Dimitris Bertsimas, Martin S Copenhaver, and Rahul Mazumder. Certifiably optimal low rank factor analysis. *The Journal of Machine Learning Research*, 18(1):907–959, 2017.
- [4] Silvère Bonnabel and Rodolphe Sepulchre. Riemannian metric and geometric mean for positive semidefinite matrices of fixed rank. *SIAM Journal on Matrix Analysis and Applications*, 31(3):1055–1070, 2010.
- [5] Silvère Bonnabel and Rodolphe Sepulchre. The geometry of low-rank Kalman filters. *Matrix Information Geometry*, pages 53–68, 2013.
- [6] Nicolas Boumal, Bamdev Mishra, P.-A. Absil, and Rodolphe Sepulchre. Manopt, a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research*, 15(1):1455–1459, jan 2014.
- [7] Claude Le Bris and Pierre Rouchon. Low rank approximation for the numerical simulation of high dimensional Lindblad and Riccati equations. *arXiv preprint arXiv:1207.4580*, 2012.
- [8] Gianluca Ceruti and Christian Lubich. Time integration of symmetric and anti-symmetric low-rank matrices and tucker tensors. *BIT Numerical Mathematics*, 60:591–614, 2020.
- [9] Valentina Ciccone, Augusto Ferrante, and Mattia Zorzi. An alternating minimization algorithm for factor analysis. *Kybernetika*, 55(4):740–754, 2019.
- [10] Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.
- [11] Geir Evensen. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res*, 99:10143–10162, 1994.
- [12] Lucia Falconi, Augusto Ferrante, and Mattia Zorzi. A robust approach to arma factor modeling. *IEEE Transactions on Automatic Control*, 2023.
- [13] H.H. Harman. *Modern Factor Analysis*. University of Chicago Press, 1967.
- [14] Uwe Helmke and Mark A Shayman. Critical points of matrix least square distance functions. In *System Structure and Control 1992*, pages 116–118. Elsevier, 1992.
- [15] Nicholas Higham. Matrix nearness problems and applications. *Applications of Matrix Theory*, pages 1–27, 03 1989.
- [16] Othmar Koch and Christian Lubich. Dynamical low-rank approximation. *SIAM Journal on Matrix Analysis and Applications*, 29:434–454, 2007.
- [17] Marc Lambert, Silvere Bonnabel, and Francis Bach. The recursive variational Gaussian approximation (R-VGA). *Statistics and Computing*, 32(1):10, 2022.

- [18] Marc Lambert, Sinho Chewi, Francis Bach, Silvère Bonnabel, and Philippe Rigollet. Variational inference via Wasserstein gradient flows. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [19] C. Le Bris and P. Rouchon. Low-rank numerical approximations for high-dimensional Lindblad equations. *Phys. Rev. A*, 87, 2013.
- [20] Mathieu Le Provost, Ricardo Baptista, Youssef Marzouk, and Jeff D Eldredge. A low-rank ensemble kalman filter for elliptic observations. *Proceedings of the Royal Society A*, 478(2266):20220182, 2022.
- [21] Wu Lin, Mohammad Khan, and Mark Schmidt. *Stein’s Lemma for the Reparameterization Trick with Exponential Family Mixtures*. 10 2019.
- [22] Christian Lubich and Ivan V Oseledets. A projector-splitting integrator for dynamical low-rank approximation. *BIT Numerical Mathematics*, 54(1):171–188, 2014.
- [23] Achim Nonnenmacher and Christian Lubich. Dynamical low-rank approximation: applications and numerical experiments. *Mathematics and Computers in Simulation*, 79(4):1346–1357, 2008.
- [24] Erkki Oja. Simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, 15(3):267–273, 1982.
- [25] Bernt Oksendal. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.
- [26] Jonathan Schmidt, Philipp Hennig, Jörg Nick, and Filip Tronarp. The rank-reduced kalman filter: Approximate dynamical-low-rank filtering in high dimensions. *arXiv preprint arXiv:2306.07774*, 2023.
- [27] Alexander Shapiro. Rank-reducibility of a symmetric matrix and sampling theory of minimum trace factor analysis. *Psychometrika*, 47:187–199, 1982.
- [28] Sebastian Thrun, Yufeng Liu, Daphne Koller, Andrew Y. Ng, Zoubin Ghahramani, and Hugh Durrant-Whyte. Simultaneous localization and mapping with sparse extended information filters. *The International Journal of Robotics Research*, 23(7-8), 2004.
- [29] Michael E. Tipping and Christopher M. Bishop. Probabilistic principal component analysis. *Journal of the Royal Statistical Society.*, 61:611–622, 1999.
- [30] Dinh Tuan Pham, Jacques Verron, and Marie Christine Roubaud. A singular evolutive extended Kalman filter for data assimilation in oceanography. *Journal of Marine Systems*, 16(3):323–340, 1998.
- [31] Bart Vandereycken. Riemannian and multilevel optimization for rank-constrained matrix problems. *Faculty of Engineering, Katholieke Universiteit Leuven*, page 13, 2010.

A Proof of Lemma 1

Letting $\delta\psi_1, \dots, \delta\psi_d$ denote the diagonal elements of $\delta\psi$, the cost (24) writes

$$\begin{aligned}\bar{C}(\delta\psi) &= \sum_{i=1}^n (h_{ii} - \delta\psi_i)^2 + 4\text{Tr}(UU^T H \delta\psi) - 2\text{Tr}(UU^T \delta\psi^2) \\ &\quad - 2\text{Tr}(U^T H U U^T \delta\psi U) + \text{Tr}(U^T \delta\psi U U^T \delta\psi U) \\ &:= \textcircled{1} + \textcircled{2} + \textcircled{3} + \textcircled{4} + \textcircled{5}.\end{aligned}$$

Let us analyze each term, and write them as functions of $\bar{\delta\psi} := (\delta\psi_1, \dots, \delta\psi_d)^T = \text{diag}(\delta\psi)$.

$$\begin{aligned}\textcircled{1} &= \sum_{i=1}^n (h_{ii} - \delta\psi_i)^2 = (\bar{h} - \bar{\delta\psi})^T (\bar{h} - \bar{\delta\psi}), \quad \bar{h} := \text{diag}(H) \\ \textcircled{2} &= 4\text{Tr}(UU^T H \delta\psi) = 4\bar{h}_U^T \bar{\delta\psi}, \quad \bar{h}_U := \text{diag}(UU^T H) \\ \textcircled{3} &= -2\text{Tr}(UU^T \delta\psi^2) = -2 \sum_{k=1}^d \delta\psi_k^2 \left(\sum_{j=1}^r U_{kj}^2 \right) = -2\bar{\delta\psi}^T \bar{D} \bar{\delta\psi}\end{aligned}$$

where we let \bar{D} be the diagonal matrix defined by $\bar{D}_{ii} = \sum_{j=1}^r U_{ij}^2$. As concerns the fourth term we have

$$\begin{aligned}\textcircled{4} &= -2\text{Tr}(U^T H U U^T \delta\psi U) \\ &= -2 \sum_{1 \leq i, j \leq r} (U^T H U)_{ij} \left(\sum_{k=1}^d \delta\psi_k U_{ki} U_{kj} \right) \\ &= -2 \sum_{k=1}^d \delta\psi_k \sum_{1 \leq i, j \leq r} (U^T H U)_{ij} U_{ki} U_{kj} \\ &= -2\Lambda^T \bar{\delta\psi}, \quad \Lambda = (\Lambda_1, \dots, \Lambda_d)^T \in \mathbb{R}^d\end{aligned}$$

with $\Lambda_k = \sum_{1 \leq i, j \leq r} (U^T H U)_{ij} U_{ki} U_{kj}$ for each k . Finally

$$\begin{aligned}\textcircled{5} &= \text{Tr}(U^T \delta\psi U U^T \delta\psi U) = \sum_{1 \leq i, j \leq r} \left(\sum_{k=1}^d \delta\psi_k U_{ki} U_{kj} \right)^2 \\ &= \sum_{1 \leq i, j \leq r} (\Gamma_{ij}^T \bar{\delta\psi})^2,\end{aligned}$$

where each $\Gamma_{ij} \in \mathbb{R}^d$ is a vector whose k -th component is $U_{ki} U_{kj}$. Thus we seek to minimize

$$\begin{aligned}\bar{C}(\bar{\delta\psi}) &= (\bar{h} - \bar{\delta\psi})^T (\bar{h} - \bar{\delta\psi}) + 4\bar{h}_U^T \bar{\delta\psi} - 2\bar{\delta\psi}^T \bar{D} \bar{\delta\psi} - 2\Lambda^T \bar{\delta\psi} \\ &\quad + \sum_{1 \leq i, j \leq r} (\Gamma_{ij}^T \bar{\delta\psi})^2.\end{aligned}$$

Thus

$$\begin{aligned}\nabla_{\bar{\delta\psi}} \bar{C}(\bar{\delta\psi}) &= 2(\bar{\delta\psi} - \bar{h}) + 4\bar{h}_U - 4\bar{D} \bar{\delta\psi} - 2\Lambda + 2 \sum_{1 \leq i, j \leq r} \Gamma_{ij} \Gamma_{ij}^T \bar{\delta\psi} \\ \Rightarrow \frac{1}{2} \nabla_{\bar{\delta\psi}} \bar{C}(\bar{\delta\psi}) &= (I - 2\bar{D} + \sum_{1 \leq i, j \leq r} \Gamma_{ij} \Gamma_{ij}^T) \bar{\delta\psi} - \Lambda + 2\bar{h}_U - \bar{h}.\end{aligned}$$

And the optimizer is given by

$$\overline{\delta\psi}^* = (I - 2\bar{D} + \sum_{1 \leq i, j \leq r} \Gamma_{ij} \Gamma_{ij}^T)^{-1} (\bar{h} - 2\bar{h}_U + \Lambda).$$

Finally, we note that $\sum_{1 \leq i, j \leq r} \Gamma_{ij} \Gamma_{ij}^T = \Upsilon \Upsilon^T$ where Υ is a $d \times p^2$ matrix whose columns consist of the p^2 vectors Γ_{ij} . To avoid double counting we may take for Υ a matrix of size $d \times \frac{p(p+1)}{2}$ whose first $\frac{p(p-1)}{2}$ columns are given by the vectors $\sqrt{2}\Gamma_{ij}$, $i < j$, and whose last r columns are the Γ_{ii} 's.

B Monitoring the quality of approximation

An advantage of the method is that one may monitor the quality of the approximation by computing the discrepancy $\|H - P(H)\|^2$ at all times, with a cost being linear in d . This opens up for methods that could dynamically adapt the parameter p , or switch between PPCA and FA. Indeed, have the following result.

Proposition 8. *In the PPCA case, we have*

$$\|H - P_{U,R,s}(H)\|^2 = \text{Tr}((I - UU^T)[H - \delta s I]^T [H - \delta s I]).$$

This may be compared to $\|H\|^2 = \|H - P(H)\|^2 + \|P(H)\|^2$ to assess the quality of the approximation.

As for the FA case, we have advantageously rewritten $\|H - P_{U,R,\psi}(H)\|^2$ as the vector cost (28), allowing for the following result.

Proposition 9. *In the FA case, $\|H - P_{U,R,\psi}(H)\|^2$ may be computed as (28) letting*

$$\alpha = \text{Tr}(H^2) - 2\text{Tr}(UU^T H^2) + \text{Tr}(UU^T HUU^T H) - \bar{h}^T \bar{h}.$$

This may be compared to $\|H\|^2 = \|H - P(H)\|^2 + \|P(H)\|^2$ to assess the quality of the approximation.

Proof. To compute the constant α in (28), we write $\tilde{C}(0) = \alpha + \bar{h}^T \bar{h}$. This is also equal to (24) at $\delta\psi = 0$, i.e., letting $\tilde{H} = H$ in the expression. This yields α . \square