

Numerical Analysis in Spaces of Fixed-Rank Tensors and Other Riemannian Manifolds

Simon Jacobsson

Dissertation presented in partial fulfillment
of the requirements for the degree of
Doctor of Engineering Science (PhD):
Computer Science

Supervisors:

Prof. dr. N. Vannieuwenhoven

Prof. dr. R. Vandebril

Prof. dr. J. Van der Veken

January 2026

NUMERICAL ANALYSIS IN SPACES OF FIXED-RANK TENSORS AND OTHER RIEMANNIAN MANIFOLDS

Simon JACOBSSON

Supervisors:

Prof. dr. N. Vannieuwenhoven

Prof. dr. R. Vandebril

Prof. dr. J. Van der Veken

Members of the

Examination Committee:

Prof. dr. ir. P. Wambacq, chair

Prof. dr. M. Zambon

Prof. dr. ir. K. Meerbergen

Prof. dr. M. Ishteva

Prof. dr. R. Zimmermann

(Syddansk Universitet)

Dissertation presented in partial
fulfillment of the requirements for
the degree of Doctor of Engineering
Science (PhD): Computer Science

January 2026

© 2025 Simon Jacobsson

Uitgegeven in eigen beheer, Simon Jacobsson, Celestijnenlaan 200A, B-3001 Leuven (Belgium)

Alle rechten voorbehouden. Niets uit deze uitgave mag worden vermenigvuldigd en/of openbaar gemaakt worden door middel van druk, fotokopie, microfilm, elektronisch of op welke andere wijze ook zonder voorafgaande schriftelijke toestemming van de uitgever.

All rights reserved. No part of the publication may be reproduced in any form by print, photoprint, microfilm, electronic or any other means without written permission from the publisher.

Preface

I want to start this thesis by thanking the people who made it possible.

First, I want to thank my three supervisors Nick, Raf, and Joeri for initiating the ManiFactor project, and then inviting me to be part of it. It has been an adventure to move to Belgium and to be part of KU Leuven. I am grateful to you Nick in particular for all the time you have spent helping me: for all the evenings and weekends where you made yourself available.

Second, I want to thank my jury members Ralf, Mariya, Karl, and Marco for agreeing to read this text. You are experts in your respective fields, and your feedback is very valuable.

Third, I want to thank my colleagues at the computer science and maths departments. In particular those I shared an office with. Astrid and Thijs, I will never forget our week in Paris. Arne, thank you for entertaining my *idée fixes*. Mateo, good luck in Sapporo and wherever you end up next. I hope it's somewhere that I can visit you. Andreas, you taught me a lot about homogeneous spaces. Nicolas, you were there as well!

Fourth, I want to thank the computer science department's system group. You went above and beyond!

Jag vill också säga tack till min Elin, för att du styr upp mig emellanåt och för att du väntade på mig.

Popularized Abstract

The *Casio fx-9750G* is an excellent multi-purpose pocket-sized scientific calculator. It can easily compute $\exp(2.121924942)$ to 10 significant digits. As a kid, I used to think it stored precomputed values of $\exp(x)$ in a table. However, storing the values for all 10^{10} possible 10-digit inputs would require 40 GB of memory, which is infeasible. Rather, it uses a *power series*

$$\exp(x) = 1 + x + \frac{1}{2}x^2 + \frac{1}{6}x^3 + \frac{1}{24}x^4 + \frac{1}{120}x^5 + \frac{1}{720}x^6 + \cdots \quad (1)$$

Truncating at the x^6 term gives an approximation,

$$\exp(x) \approx 1 + x + \frac{1}{2}x^2 + \frac{1}{6}x^3 + \frac{1}{24}x^4 + \frac{1}{120}x^5 + \frac{1}{720}x^6, \quad (2)$$

which is accurate close to $x = 0$. The right-hand side of (2) is a sequence of additions, subtractions, multiplications, and divisions. Hence the *fx-9750G* has to implement these four operations, called *basic operations*, and then it only needs to store the values $(1, 1, \frac{1}{2}, \frac{1}{6}, \frac{1}{24}, \frac{1}{120}, \frac{1}{720})$ in memory. *Numerical analysis* is the construction and study of such algorithms.

Problems that catch the attention of numerical analysts are not limited to computing exponentials and logarithms. For example they can be about equations that involve quantities and their rate of change, so-called *differential equations*, or large linear systems, called *matrix equations*. The aim is to reduce solving these to sequences of basic operations.

Numerical analysis is also concerned with how efficiently this can be done. Note that the coefficients in (1) get smaller and smaller. Their precise decay rate is important for deciding where to truncate the series. Faster decay means

we can truncate sooner. Interestingly, there is a trick to make them decay faster:

$$\begin{aligned} \exp(x) &= [\exp(x/2)]^2 \\ &= \left[1 + \frac{1}{2}x + \frac{1}{8}x^2 + \frac{1}{48}x^3 + \frac{1}{384}x^4 + \frac{1}{3840}x^5 + \frac{1}{46080}x^6 + \dots \right]^2. \quad (3) \end{aligned}$$

The tradeoff is that we have to do one multiplication more at the end, but this is only *one* basic operation extra, which is easily outweighed by not having to use as many terms.

Low-rank matrices and tensors are numerical analytic tools to compress linear and multilinear data. To store a general matrix or tensor requires more memory than to store one with low rank. Moreover, multiplying two matrices, which is a fundamental operation in linear algebra, can be done more efficiently with lower rank. The smaller the rank, the bigger the benefit. For this reason, low-rank methods are ubiquitous in numerical analysis.

In this thesis, we are interested in a special type of abstraction, called a *manifold*. These abstractions invert the central question in numerical analysis. Rather than ask “given these operations, what can I compute?”, we look at the question “given this computation, what operations do I need?”. The result is a family; *topological*, *differentiable*, *smooth*, *Riemannian*, and so on; of structures that we can put on abstract sets. If an algorithm makes use of derivatives then it only makes sense in the setting of a *differentiable manifold*, if it also uses distances between points then it needs to be in the setting of a *Riemannian manifold*, or if it only uses distances between points but not derivatives then the setting can be relaxed slightly to a *metric space*.

One part of our contribution is to the understanding of low-rank tensors as smooth and Riemannian manifolds. Another part is to the understanding of how approximation errors behave on general Riemannian manifolds.

Gepopulariseerde Samenvatting

De *Casio fx-9750G* is een uitstekende multifunctionele wetenschappelijke rekenmachine in zakformaat. Hij kan gemakkelijk $\exp(2.121924942)$ tot 10 significante cijfers berekenen. Als kind dacht ik altijd dat hij vooraf berekende waarden van $\exp(x)$ in een tabel opsloeg. Het opslaan van de waarden voor alle mogelijke 10^{10} 10-cijferige invoer zou echter 40 GB geheugen vereisen, wat onhaalbaar is. In plaats daarvan wordt een *machtrees* gebruikt

$$\exp(x) = 1 + x + \frac{1}{2}x^2 + \frac{1}{6}x^3 + \frac{1}{24}x^4 + \frac{1}{120}x^5 + \frac{1}{720}x^6 + \dots \quad (4)$$

Afbreken bij de term x^6 geeft een benadering,

$$\exp(x) \approx 1 + x + \frac{1}{2}x^2 + \frac{1}{6}x^3 + \frac{1}{24}x^4 + \frac{1}{120}x^5 + \frac{1}{720}x^6, \quad (5)$$

wat nauwkeurig is tot bijna $x = 0$. De rechterkant van (5) is een reeks optellingen, aftrekkingen, vermenigvuldigingen en delingen. De *fx-9750G* moet dus deze vier bewerkingen, de zogenaamde *basisbewerkingen*, uitvoeren en hoeft vervolgens alleen de waarden $(1, 1, \frac{1}{2}, \frac{1}{6}, \frac{1}{24}, \frac{1}{120}, \frac{1}{720})$ in het geheugen op te slaan. *Numerieke analyse* is de constructie en studie van dergelijke algoritmen.

Problemen die de aandacht van numerieke analisten trekken, beperken zich niet tot het berekenen van exponentiële en logaritmische uitdrukkingen. Ze kunnen bijvoorbeeld gaan over vergelijkingen die betrekking hebben op grootheden en hun veranderingssnelheid, zogenaamde *differentiaalvergelijkingen*, of grote lineaire systemen, zogenaamde *matrixvergelijkingen*. Het doel is om hun oplossingen te reduceren tot reeksen basisbewerkingen.

Numerieke analyse houdt zich ook bezig met hoe efficiënt dit kan worden gedaan. Merk op dat de coëfficiënten in (4) steeds kleiner worden. Hun precieze vervalsnelheid is belangrijk om te bepalen waar de reeks moet worden afgebroken. Sneller verval betekent dat we eerder kunnen afbreken. Interessant genoeg is er

een truc om ze sneller te laten vervallen:

$$\begin{aligned} \exp(x) &= [\exp(x/2)]^2 \\ &= \left[1 + \frac{1}{2}x + \frac{1}{8}x^2 + \frac{1}{48}x^3 + \frac{1}{384}x^4 + \frac{1}{3840}x^5 + \frac{1}{46080}x^6 + \dots \right]^2. \quad (6) \end{aligned}$$

Het nadeel is dat we aan het eind één vermenigvuldiging extra moeten doen, maar dit is slechts één basisbewerking extra, wat gemakkelijk teniet wordt gedaan doordat er minder termen gebruikt hoeven te worden.

Lage-rang matrices en tensoren zijn een numeriek analytisch hulpmiddel om lineaire en multilineaire data te comprimeren. Het opslaan van een algemene matrix of tensor vereist meer geheugen dan het opslaan van een matrix of tensor met een lage rang. Bovendien kan het vermenigvuldigen van twee matrices, een fundamentele bewerking in de lineaire algebra, efficiënter worden uitgevoerd met een lagere rang. Hoe kleiner de rang, hoe groter het voordeel. Om deze reden zijn lage-rang methoden alomtegenwoordig in de numerieke analyse.

In dit proefschrift zijn we geïnteresseerd in een speciaal type abstractie, een zogenaamde *variëteit*. Deze abstracties keren de centrale vraag in de numerieke analyse om. In plaats van te vragen “Wat kan ik met deze bewerkingen berekenen?”, kijken we naar de vraag “Welke bewerkingen heb ik nodig voor deze berekening?”. Het resultaat is een familie; topologische, differentieerbare, gladde, Riemannse, enzovoort; van structuren die we op abstracte verzamelingen kunnen plaatsen. Als een algoritme afgeleiden gebruikt, is het alleen zinvol in de context van een differentieerbare variëteit. Als het ook afstanden tussen punten gebruikt, moet het in de context van een Riemannse variëteit staan. Als het alleen afstanden tussen punten gebruikt, maar geen afgeleiden, kan de context enigszins worden versoepeld tot een metrische ruimte.

Een deel van onze bijdrage is het begrijpen van lage-rang tensoren als gladde en Riemannse variëteiten. Een ander onderdeel betreft het begrijpen hoe benaderingsfouten zich gedragen op algemene Riemannse variëteiten.

Abstract

We consider the geometry of low-rank tensors in two different ways. First, we look at the Lie group action induced by changes of basis in each mode. We show that this action is transitive on the set of

- tensors with canonical polyadic rank r and multilinear rank (r, \dots, r) ,
- tensors with tensor train-rank (s_1, \dots, s_{d-1}) and multilinear-rank $(s_1, s_1 s_2, \dots, s_{d-2} s_{d-1}, s_{d-1})$,
- tensors with multilinear-rank (t_1, \dots, t_d) such that $t_1 = t_2 \cdots t_d$,

whenever r, s_1, \dots, s_{d-1} are sufficiently small. This induces a homogeneous smooth manifold structure.

We explore a few natural applications of this structure. Most importantly, it induces a so-called *canonical Riemannian metric* on these sets. Geodesics are then induced by horizontal geodesics on the change-of-basis Lie group, which is a product of general linear groups. Using a known expression for geodesics on the general linear group, along with a clever choice of representative where the horizontal space has a particular rank structure, we derive an expression for geodesics on low-rank tensors. Importantly, these can then be evaluated efficiently using ideas from the theory of exponential integrators.

We further classify the manifolds that have an important additional property, namely those that are *reductive*. This property implies that the manifold is Riemannian-homogeneous, meaning that changes of basis are isometries.

Secondly, we continue the work of Lars Swijsen on the Riemannian geometry of rank-1 tensors as an embedded submanifold of Euclidean space. By identifying the metric locally as a *warped product*, we are able to derive the exponential and logarithmic maps. We are moreover able to put this manifold in a more general context. There is a one-parameter family of metrics, obtained from tweaking

a parameter in the *warping function* of the warped product. It includes the Euclidean metric, but also some other interesting geometries. In particular, some of these are *geodesically connected*, meaning that any pair of points can be connected by a geodesic. This property is important in for example *consensus aggregation* where we want to find the rank-1 tensor that best represents the average of a set of many rank-1 tensors.

We then consider the general problem of how to approximate a map into a Riemannian manifold. Many established approximation schemes follow a general pullback-approximate-pushforward template. We relate their accuracy to the geometry of the target manifold by showing that if one uses the exponential and logarithmic maps to pull back and push forward to and from the tangent space, then it is possible to bound the approximation error on the manifold in terms of the linear approximation on the tangent space and a lower bound on the sectional curvature. It turns out that for many manifolds that appear organically in numerical analysis, we have explicit expressions for the exponential and logarithmic maps, and there are often global lower curvature bounds as well. Numerical experiments and a natural application to *Krylov subspace recycling* are presented.

Our analysis uses *Toponogov's theorem*, a classic result from Riemannian comparison theory about geodesic triangles. The idea is that one can construct such a triangle where one corner is the true function value and another corner is the approximated value. Having constructed the correct *comparison triangle*, the analysis is straightforward.

If the exponential and logarithmic maps are not known for a manifold, one can use a so-called *retraction*, which is a map that agrees with the exponential to first order. We discuss how it is possible to use our framework to think about the approximation error when retracting. Typically, this will overestimate the error. Numerical experiments confirm that there is practically no difference in the errors between exponentiating and retracting.

We briefly consider the related problem of integrating an ordinary differential equation on a manifold. Importantly, the central problem in *dynamical low-rank approximation* is to integrate an ODE on a low-rank manifold. We construct two new types of dynamical low-rank integrators based on our insights into the geometry of low-rank tensors. One type uses the Riemannian structure and is the so-called *geodesic Euler* methods, which are generalizations of Euler forward and backward that use geodesics to step on the manifold. The other type combines the facts that the Lie exponential is always a retraction and that retractions on a Lie group induce retractions on its quotients. We note that this second type is not the same as *exponential integrators* because it depends

on the precise form of the horizontal space. While exponential integrators often use Krylov subspace methods for efficient computations, our integrator thus is able to use *low-rank* methods instead.

Everything numerical is implemented in Julia, and the source code is available. The low-rank tensor geometries and the rank-1 warped product geometries are all contributed to the popular library `Manifolds.jl`. There is also a separate package, `ManiFactor.jl`, published to demo the process of approximating maps into Riemannian manifolds via this pullback-approximate-pushforward technique.

Beknopte samenvatting

We beschouwen de meetkunde van laagrangtensoren op twee verschillende manieren. Ten eerste bekijken we de Liegroepactie die wordt veroorzaakt door basisveranderingen in elke modus. We laten zien dat deze actie transitief is op de verzameling van

- tensoren met canonieke polyadische rang r en multilineaire rang (r, \dots, r) ,
- tensoren met tensor trein-rang (s_1, \dots, s_{d-1}) en multilineaire rang $(s_1, s_1 s_2, \dots, s_{d-2} s_{d-1}, s_{d-1})$,
- en tensoren met multilineaire rang (t_1, \dots, t_d) zodat $t_1 = t_2 \cdots t_d$,

wanneer r, s_1, \dots, s_{d-1} zijn voldoende klein.

Dit induceert een homogene gladde variëteitstructuur.

We onderzoeken een paar natuurlijke toepassingen van deze structuur. De belangrijkste is dat ze een zogenaamde *canonieke Riemannse metriek* induceert. Geodeten worden geïnduceerd door horizontale geodeten op de Liegroep van de basisverandering, die een product is van algemene lineaire groepen. Met behulp van een bekende uitdrukking voor geodeten op de algemene lineaire groep, samen met een slimme keuze van representatieve waarden waarbij de horizontale ruimte een bepaalde rangstructuur heeft, leiden we een uitdrukking af voor geodeten op lagerangtensoren. Belangrijk is dat deze vervolgens efficiënt kunnen worden geëvalueerd met behulp van ideeën uit de theorie van exponentiële integratoren.

We classificeren deze variëteiten wanneer ze een belangrijke extra eigenschap hebben, namelijk wanneer ze *reductief* zijn. Deze eigenschap impliceert dat de Riemannse variëteit homogeen is, wat betekent dat basisveranderingen isometrieën zijn, en heeft implicaties voor zogenaamde *Lieintegratoren*.

Ten tweede zetten we het werk van Lars Swiisen aan de Riemannse meetkunde van rang-1-tensoren als een ingebedde deelvariëteit van de Euclidische

ruimte voort. Door de metriek lokaal te identificeren als een *warped product*, kunnen we de exponentiële en logaritmische afbeeldingen afleiden. Bovendien kunnen we deze variëteit in een meer algemene context plaatsen. Er is een éénparameterfamilie van metrieken, verkregen door een parameter in de *warpingfunctie* van het warped product aan te passen. Deze omvat de Euclidische, maar ook enkele andere interessante geometrieën. Sommige hiervan zijn met name *geodesisch verbonden*, wat betekent dat elk paar punten door een geodeet met elkaar verbonden kan worden. Deze eigenschap is bijvoorbeeld belangrijk bij *consensusaggregatie*, waarbij we de rang-1-tensor willen vinden die het beste het gemiddelde van een verzameling van vele rang-1-tensoren vertegenwoordigt.

Vervolgens bekijken we het algemene probleem van hoe we een afbeelding in een Riemannse variëteit kunnen benaderen. Veel gevestigde benaderingsschema's volgen een algemeen pullback-benadering-pushforward-sjabloon. We relateren hun nauwkeurigheid aan de meetkunde van de doelvariëteit door aan te tonen dat als men de exponentiële en logaritmische afbeeldingen gebruikt om terug te trekken en vooruit te duwen van en naar de raakruimte, het mogelijk is om de benaderingsfout op de variëteit te begrenzen in termen van de lineaire benadering op de raakruimte en een ondergrens voor de sectionele kromming. Het blijkt dat we voor veel variëteiten die organisch voorkomen in numerieke analyse expliciete uitdrukkingen hebben voor de exponentiële en logaritmische afbeeldingen, en dat er vaak ook globale ondergrenzen voor de kromming zijn. Numerieke experimenten en een natuurlijke toepassing op *Krylov-deelruimte recycling* worden gepresenteerd.

Onze analyse maakt gebruik van de *stelling van Toponogov*, een klassiek resultaat uit de Riemannse vergelijkingstheorie over geodetische driehoeken. Het idee is dat men zo'n driehoek kan construeren waarbij één hoekpunt de exacte functiewaarde is en een ander hoekpunt de benaderende waarde. Nadat de juiste *vergelijkingsdriehoek* is geconstrueerd, is de analyse eenvoudig.

Als de exponentiële en logaritmische afbeeldingen voor een variëteit niet bekend zijn, kan men een zogenaamde *retractie* gebruiken, een afbeelding die overeenkomt met de exponentiële waarde tot op de eerste orde. We bespreken hoe het mogelijk is om ons raamwerk te gebruiken om na te denken over de benaderingsfout bij het retractie. Meestal zal dit de fout overschatten. Numerieke experimenten bevestigen dat er geen praktisch verschil is in de fouten tussen exponentiëren en retractie.

We bespreken kort het gerelateerde probleem van het integreren van een gewone differentiaalvergelijking op een variëteit. Het centrale probleem bij dynamische benaderingen van lage rang is het integreren van zo'n vergelijking

op een variëteit van lagerangtensoren. We construeren twee nieuwe types dynamische integratoren van lage rang op basis van onze inzichten in de meetkunde van tensoren van lage rang. Het ene type maakt gebruik van de Riemannse structuur en betreft de zogenaamde geodetische Euler-methoden. Dit zijn veralgemeningen van de voorwaartse en achterwaartse Eulermethode die geodeten gebruiken om op de variëteit te bewegen. Het andere type combineert de feiten dat de Lie-exponentiële altijd een retractie is en dat retracties op een Liegroep retracties op zijn quotiënten induceren. We merken op dat dit tweede type niet hetzelfde is als exponentiële integratoren, omdat het afhangt van de precieze vorm van de horizontale ruimte. Hoewel exponentiële integratoren vaak Krylov-deelruimte gebruiken, terwijl exponentiële integratoren vaak Krylovdeelruimten gebruiken, kan onze integrator in plaats daarvan lagerangmethodes gebruiken.

We hebben alles numeriek geïmplementeerd in Julia en de broncodes beschikbaar gemaakt. De meetkundes van de lagerangtensoren en de warped producten van rang-1-tensoren zijn toegevoegd aan de populaire bibliotheek `Manifolds.jl`. We hebben ook een afzonderlijk pakket, `ManiFactor.jl`, gepubliceerd om het benaderen van afbeeldingen naar Riemannse variëteiten via de pullback-benadering-pushforward-techniek te demonstreren.

Contents

Popularized Abstract	iii
Gepopulariseerde Samenvatting	v
Abstract	vii
Beknopte samenvatting	xi
List of Abbreviations	xv
List of Symbols	xv
Contents	xv
List of Figures	xvii
1 Introduction	1
1.1 Manifold methods	1
1.2 Low-rank methods	6
1.3 Research questions and objectives	8
1.4 Approach and research methods	8
1.5 Outline	8
2 Preliminaries	11
2.1 Smooth manifolds	11
2.2 Riemannian manifolds	15
2.3 Lie groups and homogeneous spaces	21
2.4 Algebraic geometry	25
2.5 Tensors	26
2.6 Appendix: proof of proposition 2.19	40
3 Low-rank tensor manifolds	41

3.1	Introduction	41
3.2	The CP manifold	43
3.3	The Tucker manifold	54
3.4	The tensor train manifold	57
3.5	Application: low-rank tensor approximation	61
3.6	Conclusion and outlook	65
3.7	Appendix: error bound for Padé approximant	66
3.8	Appendix: counting the operations	67
4	Segre–Veronese manifolds	71
4.1	Introduction	71
4.2	α -warped geometries	72
4.3	The exponential map	75
4.4	The logarithmic map	79
4.5	Sectional curvature	88
4.6	Application: consensus aggregation	90
4.7	Conclusion and outlook	92
4.8	Appendix: proofs of technical lemmas	93
5	Approximation theory on manifolds	97
5.1	Introduction	97
5.2	Classical approximation theory	100
5.3	Approximation theory on manifolds	106
5.4	Manifolds with bounded sectional curvature	114
5.5	A concrete algorithm	124
5.6	Application: Krylov subspaces	126
5.7	Application: dynamical low-rank approximation	128
5.8	Effects of using a retraction	129
5.9	Conclusion and outlook	131
6	Julia implementations	133
6.1	Tests	134
7	Conclusion	137
7.1	Future work	139
	Bibliography	141

List of Figures

2.1	Proposition 2.11 gives conditions for which the length of the opposite (dashed) side in M is bounded by the length of the opposite (dashed) side in N . Equal-length geodesics and equal angles have the same number of notches.	19
2.2	Proposition 2.12 gives conditions for which the length of the sector (dashed) in M is bounded by the length of the sector (dashed) in N . Equal-length geodesics and equal angles have the same number of notches/lines.	20
2.3	A mathematical expression.	28
2.4	Composition of linear maps.	29
2.5	Penrose diagram for the CP decomposition. Edges are labeled with the vector space they represent, and $V_i = [v_i^1 \ \dots \ v_i^r]$. I is the diagonal tensor $I_{i_1 \dots i_d} = 1$ if $i_1 = \dots = i_d$ else $I_{i_1 \dots i_d} = 0$	30
2.6	Penrose diagram for the Tucker decomposition (2.46). Edges are labeled with the vector space they represent.	35
2.7	Penrose diagram for the TT decomposition (2.51). Edges are labeled with the vector space they represent.	36
3.1	Errors for different low-rank approximations of (3.78).	63
3.2	Errors for different low-rank approximations of (3.79).	64
4.1	Geodesics between $(0, 1)$ and $(1, 0)$ in the α -warped geometry for $\alpha = 0.01$ (the red path approaching the semicircle), $\frac{1}{4}$, $\frac{1}{2}$, $\frac{3}{4}$, 1 (the purple straight line), $\frac{5}{4}$, $\frac{3}{2}$, $\frac{7}{4}$, and 1.99 (the blue path that narrowly avoids the puncture at the origin).	73
4.2	We recreate the setup from [CST23, Section 4.1] and compare the Riemannian distance between the approximate decompositions and the ground truth.	92

5.1	Using different number of samples, we approximate $f: [-1, 1]^2 \rightarrow S^2$. Let Γ be a 10×10 grid on $[-1, 1]^2$. Solid lines show $f(\Gamma)$ and dashed lines show $\widehat{f}(\Gamma)$	99
5.2	Geodesic triangles defined by p , f , and \widehat{f} . The map from M to N is $\exp_{N,q} \circ \log_{M,p}$	107
5.3	The fattest possible comparison triangle is still contained in one hemisphere $T = B_\sigma(q)$	113
5.4	γ is a horizontal geodesic in G emanating from 1. It cuts the fiber gH at a right angle in g	119
5.5	Approximation error for the map in (5.91) into the Grassmann manifold $\text{Gr}(n, k)$, compared against what is predicted by theorem 5.5, for different numbers of Chebyshev nodes N	127
5.6	Approximation error for the map in (5.93) into the Segre manifold $\text{Seg}(\mathbb{R}^n \times \mathbb{R}^n)$, compared against what is predicted by theorem 5.5, for different numbers of Chebyshev nodes N in each variable.	129
5.7	Approximation error for the map in (5.91), using different retraction methods.	129
5.8	Time it takes to evaluate $\widehat{f}(x)$, using different retraction methods.	130

1. *Introduction*

1.1 **Manifold methods**

Manifold methods are becoming increasingly ubiquitous in numerical analysis and numerical linear algebra. Here are some examples of sets that appear in applications which are also manifolds:

- orthogonal matrices appear in motion tracking, matrix decompositions, differential equations, statistics, and optimization [AMS08; BGW15; CV19; Zim20; ZB24];
- symmetric positive definite matrices and tensors appear in diffusion tensor imaging, materials modeling, computer vision, and statistics [Moa06; Che+16; Pen20; PFA06; TP21; TP23];
- fixed-rank matrices and tensors appear in compression, tensor completion, and dynamic low-rank approximation [KL07; OT10; Ose11; VAV12; SvV22];
- fixed-dimensional linear subspaces appear in Krylov subspace methods and subspace tracking [Yan95; BGW15; SSK20; ZH22];
- sequences of nested linear subspaces appear in principal component analysis [BP23; MCB24].

Informally, a *manifold* is a set where notions from calculus are well-defined. Given a map from the set into the reals, we can ask questions like whether this map is smooth, or what are its stationary points? For a numerical analyst, a course in differential geometry answers the question “what structures a set needs for us to implement Newton’s method on it?”. We now explain how some manifolds appear organically in applications.

Problem 1. Find the smallest eigenvalue of an $n \times n$ symmetric matrix A .

To tackle this, we define the *Rayleigh quotient*

$$R(x) = \frac{x^\top Ax}{x^\top x}. \quad (1.1)$$

By the following argument, problem 1 is equivalent to minimizing R over the manifold S^{n-1} of unit length vectors in \mathbb{R}^n . First, $R(x)$ has the same value if we rescale x . It is thus indeed a function from S^{n-1} . Second, let v_1, \dots, v_n be an orthogonal eigenbasis to A with eigenvalues $\lambda_1 \leq \dots \leq \lambda_n$, and let $x = c_1 v_1 + \dots + c_n v_n$ be unit length. Then

$$R(x) = c_1^2 \lambda_1 + \dots + c_n^2 \lambda_n. \quad (1.2)$$

From here, we see that R is minimized by setting $c_1 = 1$ and $c_2 = \dots = c_n = 0$, that is, setting x equal to the eigenvector with the smallest eigenvalue. We therefore can write problem 1 equivalently as the following.

Problem 2. Minimize R over S^{n-1} .

Importantly, what started out as a numerical linear algebraic problem is now a manifold optimization problem.

Another example of manifolds appearing organically is in *Krylov subspace recycling*. Consider the following problem.

Problem 3. Solve the sequence $A_i x = b$ of sparse matrix equations efficiently.

Krylov subspace methods approach this by looking for a solution in a *Krylov subspace* $\mathcal{K}_k(A_i, b)$, spanned by vectors $b, A_i b, \dots, A_i^{k-1} b$. The idea is that the dimension k of the Krylov subspace can in practice often be chosen to be small compared to the size of the problem, and hence x can be estimated efficiently. Moreover, the elements A_i in the sequence are often similar, and we expect their Krylov subspaces to be similar as well. In Krylov subspace recycling, we therefore try to approximate A_i 's subspace, given information about nearby subspaces.

As mentioned, the set of k -dimensional linear subspaces of \mathbb{R}^n form a manifold $\text{Gr}(n, k)$, called a *Grassmannian*. Solving problem 3 using Krylov subspace recycling is thus the same as solving the following problem.

Problem 4. Approximate the sequence $\mathcal{K}_k(A_i, b)$ of points on $\text{Gr}(n, k)$.

Similarly to the previous example, what started out as a numerical linear algebraic problem is now a manifold approximation problem.

Another area where the manifold perspective is useful is *statistics*. In Euclidean space \mathbb{R}^n , the *mean* of a set x_1, \dots, x_N of points can be defined as

$$\bar{x} = \frac{1}{N}(x_1 + \dots + x_N), \quad (1.3)$$

or equivalently as

$$\bar{x} = \operatorname{argmin}_{x \in \mathbb{R}^n} ((x_1 - x)^2 + \dots + (x_N - x)^2). \quad (1.4)$$

Since there is no intrinsic addition operation on a general manifold, it is not obvious how to generalize means to manifolds from the first definition. The second definition, while slightly convoluted, is easier to generalize. The *Karcher mean* [Kar77] of a set p_1, \dots, p_N of points on a Riemannian manifold M is

$$\bar{p} = \operatorname{argmin}_{p \in M} (d(p_1, p)^2 + \dots + d(p_N, p)^2). \quad (1.5)$$

Here, d is the Riemannian distance function.

Such a mean is very useful because, similarly to the Euclidean mean, it represents the center of a set of data points. Imagine for example that we measure the orientation of some object many times, and the measurements are stored in rotation matrices. What is a good guess for the true rotation? Taking the entrywise mean of the rotation matrices will likely not yield another rotation matrix. For example, the mean of $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$ and $\begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}$ is $\begin{bmatrix} \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 1 & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} \end{bmatrix}$, which is not even invertible. Rather, something like the Karcher mean has to be used.

The failure of the sum to be orthogonal captures something fundamental about the set of rotation matrices, and about manifolds in general. In fact there is *no* sensible definition of addition on the set of 3D rotation matrices. Any definition will fail some important property like commutativity, continuity, nontriviality, etc. For more general manifolds, the existence of an addition operation implies that the manifold is equivalent to Euclidean space in a sense that we will make precise in chapter 2.

Yet another active research area, of special interest in our thesis, is the approximation of maps into manifolds. While the approximation theory of maps

into vector spaces is a mature subject, the interest in approximating maps into manifolds is still very recent.

One concrete approach to approximate maps into manifolds is via *subdivision schemes*. A Euclidean subdivision scheme is an algorithm to recursively partition a mesh on some Euclidean domain into finer meshes, typically by dividing selected elements of the mesh “in half”. The idea is then to approximate the function as a weighted sum of its values in nearby points. To build a manifold subdivision scheme, one idea is to replace the Euclidean mean with a Karcher mean. Manifold subdivision schemes have been investigated by for example [Moo16; Moo17; DS17a; DS17b].

Another approach to approximate maps into manifolds is via *curve fitting*. The Euclidean cubic spline is optimal in the sense that it minimizes the functional

$$F(S) = \int_0^1 S''(t)^2 dt, \quad S: [0, 1] \rightarrow \mathbb{R}^n, \quad S(0) = a, S(1) = b, \quad (1.6)$$

where a and b are fixed. The natural manifold generalization of this uses the *inner product* that comes with the tangent space at every point on a Riemannian manifold. A *Riemannian cubic spline* is therefore something that minimizes the functional

$$F(S) = \int \langle S''(t), S''(t) \rangle_{S(t)} dt, \quad S: [0, 1] \rightarrow M, \quad S(0) = a, S(1) = b. \quad (1.7)$$

Such approximation schemes have been explored in for example [Sam+12; BG18; HRW19; ZN19; GMA19; HW21; Wan+25].

The last approach that we’ll mention is the *pullback-approximate-pushforward* approach. It is in some sense the most versatile, but it requires knowledge of some concrete local parametrization of the manifold around each point. If $\phi: \mathbb{R}^n \rightarrow M$ is such a parametrization, and f is a map into M , then $g = \phi^{-1} \circ f$ is a map into \mathbb{R}^n , and so it can be approximated using traditional methods. This approximation, \hat{g} , can then in turn be pushed forward to the manifold by defining

$$\hat{f} = \phi \circ \hat{g}. \quad (1.8)$$

This approach has been investigated by for example [BG18; Zim20; ZH22; ZB24; SCK24; BDS24; JZ25].

Common to these approximation schemes is that they are defined intrinsically on the manifold. We mention that there are also approaches [Moa02; GS13; GL18; HL23] that embed the manifold in a Euclidean space, and perform classical approximation there, but that these are generally not state of the art.

Problems typically arise with 1) conditioning, exemplified by taking the average of two almost antipodal points on an embedded sphere; or 2) computational complexity, which is a big problem in cases like low-rank matrices where the natural embedding has very large dimension.

While the errors introduced by approximation algorithms are well-understood in the classical Euclidean setting, we sometimes lack an understanding in the more general Riemannian setting. One of our goals in this thesis is hence to find a framework where it is possible to do such an error analysis. Our framework should be broad enough that many of the aforementioned works are included, but specific enough that errors can meaningfully be analyzed. We settle on a class of Riemannian manifolds with lower bounded sectional curvature, where the approximation can be viewed as a pushforward of an approximation in the tangent space.

The condition that the sectional curvature be lower bounded comes from *Toponogov's theorem*, a comparison theorem about geodesic triangles that we use. Assuming this condition, we can construct a geodesic triangle on the target manifold such that one of its side lengths corresponds to the approximation error. Applying Toponogov's theorem, we are thus able to bound this error in terms of the error in the tangent space, which can be interpreted as the approximation error for solving a corresponding linearized problem. Previous works [Zim20, theorem 3.1] [DCN25, theorem 3.4] in this direction have used comparison theorems for the *Jacobi fields*, which only encode local information. As a result, these works are only able to derive first-order error bounds. The strength of using Toponogov's theorem is that we can derive *exact* error bounds.

Given a smooth manifold, we have some freedom in how to choose a Riemannian structure. Consider for example the set of $n \times n$ rank- r matrices. We can define its Riemannian structure in two inequivalent ways. First, as a submanifold of all $n \times n$ matrices, equipped with the Riemannian structure that it inherits from there. Second, multiplying from the left and right with invertible matrices preserves the rank, and compatibility with this action induces a Riemannian structure. We might say that the first of these geometries better reflects how we think about the set of rank- r matrices. However, between the two, the second geometry has some desirable properties: we have explicit formulas for its *geodesics*, and its curvature is bounded both from above and below. The first geometry satisfies neither of these; see [FL18, theorem 24] and [AO15, section 3.10].

Thus another goal is to review what is known about different geometries of the common matrix and tensor manifolds, so that we can relate this to our error analysis.

1.2 Low-rank methods

Low-rank methods are ubiquitous in numerical analysis and numerical linear algebra. One reason is that it takes fewer numbers to specify a low-rank matrix or tensor than to specify one with full rank. Moreover, many things we want to do to matrices and tensors; like multiplying, inverting, or factorizing; can be done more efficiently with lower rank. The main reason however is that low-rank matrices and tensors are themselves ubiquitous. From big data to approximation theory, matrices and tensors are *observed* to be, numerically or exactly, low rank.

For matrices, there is a single notion of *rank*. This is essentially the statement of the *rank-nullity theorem*. On the other hand, there are several inequivalent notions of *tensor rank*. One notion is to say that the rank of a tensor is the minimum number of rank-1 tensors that has to be summed to reach it, and to define a rank-1 tensor as something in the image of the outer product. This leads to *canonical polyadic* rank. Another notion is to list the matrix ranks of all possible single-mode unfoldings. This leads to *multilinear* or *Tucker* rank. Importantly, a tensor's canonical polyadic rank does not determine its multilinear rank, or vice versa. In this thesis we also consider the *Tensor train* rank, whose definition we discuss later. There are many other interesting notions of rank, but these are probably the three most used.

For any notion of tensor rank, there is a corresponding notion of *decomposition*. There is the canonical polyadic decomposition, the Tucker decomposition or higher order SVD, and the tensor train decomposition. The computation and approximation of these decompositions is the main research question in numerical multilinear algebra [KB09; ES09; Ish+09; SL10; CJ10; OT10; ADK11; Lan12; HRS12b; VVM12; RU13; PTC13; SVD13; Sav14; KSV14; Cic+15; Sai16; ORU18; BV18; ORU23; VVD25]. Approaches range from the optimization formulations that we mentioned in section 1.1 to successive matrix decompositions applied to each unfolding. We mention also that using randomized matrix decompositions in this second setting is an important tool to achieve competitive decomposition algorithms [HMT11; Osi25; CK25].

One active application area of low-rank methods is multivariate function approximation. The idea behind this is to start with the observation that we can store the values of a multivariate function in a multidimensional array, and then interpret this array as a tensor. Say we evaluate a function $g: [0, 1]^d \rightarrow \mathbb{R}$ on N equispaced points in each direction. This yields a d -dimensional array with N^d real numbers. Rather than compute this whole array, we can compute a low-rank approximation to it. There are two main approaches to do this.

The first option is to make the low-rank assumption on the functional level: so that multivariate functions are approximated as finite sums of products of univariate functions. The term *separable* is used in some contexts. Then the problem is discretized after this assumption has been made. The second option is to discretize first, and then assume that the resulting matrix or tensor is low-rank. Some recent investigations into this are [BGM09; GKT13; Ose13; BEM16; HT17; GKM19; DKS21; SSK24; SK23].

Each approach has its respective advantages and disadvantages. In particular for the error analysis [SSK24; GH23] of these algorithms, the approaches are quite different. Similarly to univariate functions, errors depend on the smoothness of the approximee. However, naively tensorizing those bounds is often not enough to fully explain the errors [RTW22; UT19].

Another application area is *recommender systems* [FO16; FO17; SvV22; FO23; SOF23]. The *MovieLens* dataset contains 32 million ratings given by 200 948 users to 87 585 movies. It is well-approximated by a low-rank tensor because tastes in movies are highly correlated. If I like *Shrek* and *Shrek 2*, then I will with high probability like *Shrek Forever After* as well. Here, low numerical rank is explained by *covariance*.

A third application area is *dynamical low-rank approximation* [KL07; KL10; LO14; Lub+13; K LW16; CC22; CKL22; HNS23; SCK24]. The idea is to project a matrix or tensor ordinary differential equation (ODE) to a low-rank manifold, and then to numerically integrate it there, intrinsically on the manifold. It can be shown theoretically, and confirmed in practice, that such a scheme is quasi-optimal in the sense that the solution to the projected ODE stays close to the projection of the solution to the ODE. A big application for this is to integrate discretized partial differential equations (PDEs).

These are only a selection of the application areas of low-rank methods. For a broader overview, we refer to the books by Comon and Jutten [CJ10], by Landsberg [Lan12], and by Hackbusch [Hac19]; the review paper by Kolda and Bader [KB09], and the review by Sidiropoulos et al. [Sid+17]. Nevertheless, we hope that they are enough to motivate another of our goals in this thesis: to explore the geometry of low-rank tensors.

As we mentioned in section 1.1, fixed-rank matrices are a manifold. The precise conditions under which fixed-rank tensors are a manifold are somewhat more complicated, and depend on what exact notion of rank we use. We elucidate this partly by introducing a natural Lie group action, under which sets of low-rank tensors become a smooth homogeneous manifold. From this identification follows several natural applications, the two most important being a Riemannian metric with efficiently computable geodesics.

Another approach we take is to view the set of rank-1 tensors as an embedded submanifold of Euclidean space. By identifying this set as a *warped product manifold*, we are able to solve for its exponential and logarithmic maps, as well as for a one-parameter generalization of its metric.

1.3 Research questions and objectives

We briefly recap the goals that we mentioned in sections 1.1 and 1.2. There are two main ones.

- Explore the geometry of low-rank tensors.
- Analyze the error of approximation schemes for maps into manifolds, using methods that are intrinsic to those manifolds. We want to do this analysis in a sufficiently general setting so that it covers previous works on approximating maps into manifolds.

1.4 Approach and research methods

Our methods come mostly from Riemannian geometry. In chapter 3, we make use of Lie group theory and homogeneous spaces, as well as Kruskal's theorem about the uniqueness of the CP decomposition. Chapter 4 mostly relies on the theory of warped products. The main tool in chapter 5 is Toponogov's theorem.

We also implement the exponential and logarithmic maps that we derive in Julia. This allows us to do a host of numerical experiments, to confirm predictions made from our theorems and to demo potential applications.

1.5 Outline

Chapter 2 recalls the necessary preliminaries for manifolds and tensors. We cover topics from the definition of smooth manifold to Klingenberg's theorem about the injectivity radius for compact manifolds, as well as from the characteristic property of the tensor product to the implementation of fast tensor decomposition methods.

Chapter 3 defines a natural Lie group action on sets of low-rank tensors, and shows that this action induces a smooth homogeneous structure on those

sets. This moreover induces a Riemannian structure, whose exponential map is described and shown how to evaluate efficiently. This chapter deals with the first research objective from section 1.3. The main novel results are theorems 3.3, 3.8, and 3.13. These statements can be seen as tensor generalizations of known matrix results. They rely on fairly well-known uniqueness results about tensor decompositions. These uniqueness results have however, to my knowledge, not previously been used in this way: to solve for homogeneous structures.

Chapter 4 considers the set of rank-1 tensors as an embedded submanifold of Euclidean space. The exponential and logarithmic maps are derived. We moreover discuss a one-parameter generalization of the Riemannian metric, with applications to consensus aggregation. This chapter also deals with the first research objective from section 1.3. The main novel results are theorems 4.7 and 4.11 about the logarithmic map on the Segre manifold, and proposition 4.15 about for what parameters the Segre manifold is geodesically connected. The exponential map for the Segre manifold was already derived in [SvV22], and our proposition 4.3 is a slight generalization.

Chapter 5 presents a framework for analyzing approximation errors intrinsically on Riemannian manifolds. A practical algorithm is discussed and implemented, along with a literature review of manifolds that fit into this framework. This chapter deals with the second research objective from section 1.3. The main novel result in this chapter is theorem 5.5 about an exact error bound on Riemannian manifolds whose sectional curvature can be lower bounded. The statement is a direct application of Toponogov's theorem from Riemannian comparison theory. The relation between curvature and error has been investigated before in different contexts, but this is, to my knowledge, the first time that the bound in Toponogov's theorem has been interpreted as an approximation error.

Chapter 6 discusses the use and implementation details of the Julia code that we wrote for this thesis.

The main chapters, Chapters 3 to 5, can all be read independently. Chapter 6 relies on all the main chapters.

2. Preliminaries

Manifolds are to geometry what linear spaces are to linear algebra. In this chapter, we recall basic notions like smooth and Riemannian manifolds. We also introduce some results from a branch of Riemannian geometry called comparison theory. Then we recall some standard tools from Lie group theory like group actions, stabilizers, and orbits. We then briefly recall standard notions from algebraic geometry before talking about tensors and tensor decompositions.

2.1 Smooth manifolds

The main reference for this section is Lee's [Lee13] book on smooth manifolds.

Euclidean spaces (\mathbb{R}^n), n -spheres (S^n), the set of $n \times n$ orthogonal matrices ($O(n)$), and the set of $n \times n$ symmetric rank- r matrices ($\text{SPSD}(r, n)$) are all examples of manifolds, but also more abstract spaces like the set of k -dimensional linear subspaces of \mathbb{R}^n ($\text{Gr}(k, n) = O(n)/(O(k) \times O(n-k))$), or the set of inner products on \mathbb{R}^n ($\text{SPD}(n) = \text{GL}(n)/O(n)$) are examples of manifolds. Informally, a smooth manifold is a space where you can do calculus. The way to make this rigorous is via *diffeomorphisms*, which are bijective smooth maps with smooth inverses.

Definition 2.1. A *smooth n -dimensional manifold* (or *smooth n -manifold*) M is a Hausdorff space together with an open cover $\mathcal{A} = \{U_\alpha\}_{\alpha=1}^\infty$ and homeomorphisms $\phi_\alpha: U_\alpha \rightarrow V_\alpha$, where the V_α are open subsets of \mathbb{R}^n , such that, for each α and β ,

$$\phi_\alpha \circ \phi_\beta^{-1}: \phi_\beta(U_\alpha \cap U_\beta) \rightarrow \phi_\alpha(U_\alpha \cap U_\beta) \quad (2.1)$$

is a diffeomorphism.

Remark. U_α and U_β can have an empty intersection, in which case (2.1) is a diffeomorphism by tautology.

We say that M is equipped with a *smooth structure*. We will not break down definition 2.1 much further, but just comment on some of the terms we used.

- A *topological space* is a set together with a notion of which subsets are open/closed. This is the structure we need in order to talk about *continuous* maps.
- A *Hausdorff space* is a topological space where any two distinct elements are contained in two disjoint open subsets. In some sense it is a “weak” condition, and most topologies that one encounters in nature are Hausdorff. A notable exception is \mathbb{R}^n with the Zariski topology.
- A *homeomorphism* is a continuous map with a continuous inverse.
- The open cover \mathcal{A} is called a *smooth atlas*. Without loss of generality, we can assume that \mathcal{A} is *maximal*, meaning that there is no open supercover that is also a smooth atlas. The benefit is that two smooth manifolds with the same underlying topological space are equivalent, meaning the identity map is a diffeomorphism, if and only if they have the same maximal smooth atlases.
- The homeomorphisms ϕ_α are called *charts* and the maps $\phi_\alpha \circ \phi_\beta^{-1}$ are called a *transition functions*. Note that transition functions map between subsets of \mathbb{R}^n , and so it makes sense to ask them to be a diffeomorphisms. What we might naively want to do is to ask the charts directly to be diffeomorphisms, but this is not well-defined a priori.

For example, the *orthogonal group* $O(n)$, comprising orthogonal $n \times n$ matrices, is a smooth manifold. When $n = 2$, any such matrix can be written on the form

$$f(t) = \begin{bmatrix} \cos t & -\sin t \\ \sin t & \cos t \end{bmatrix}. \quad (2.2)$$

Note that f is not invertible on \mathbb{R} . The charts $\phi_1 = f|_{(0,2\pi)}^{-1}$, $\phi_2 = f|_{(\pi,3\pi)}^{-1}$ make $O(2)$ a smooth manifold.

A more tautological example is the set of $m \times n$ matrices with full rank, denoted $\mathbb{R}_*^{m \times n}$. It is covered with one chart, because it is an open subset of \mathbb{R}^{mn} . Similarly, \mathbb{R}_*^n is just punctured Euclidean space.

From here we can define smooth maps between manifolds. A map $f: M \rightarrow N$ between smooth manifolds is *smooth* if $\phi_\alpha \circ f \circ \phi_\beta$ is smooth for each pair of charts. Importantly, this allows us to ask whether M and N are *diffeomorphic*, that is if there exists a diffeomorphism between them. This is the natural equivalence relation between smooth manifolds.

In real analysis, the *derivative* of a map $f: A \rightarrow B$ between linear spaces at $x \in A$ is the linear map $df_x: A \rightarrow B$ such that $(f - f(x) - df_x)(y)$ vanishes $o(y)$. However, a manifold is not (necessarily) a linear space. In order to define derivatives on manifolds, we will first need to associate, to each point, a linear space where the derivative can then live.

Definition 2.2. Let M be a smooth manifold and let $p \in M$. A *tangent vector at p* is a smooth curve through p , i.e. a smooth map $\gamma: \mathbb{R} \rightarrow M$ such that $\gamma(0) = p$. Let $\phi: U \rightarrow V$ be a chart around p , i.e. $p \in U$. Two tangent vectors γ and λ are *equal* if their derivatives $(\phi \circ \gamma)'(0)$ and $(\phi \circ \lambda)'(0)$ are equal. The *tangent space* $T_p M$ is the set of possible tangent vectors at p .

One readily verifies that the tangent space is a linear space, and that having the same tangent vectors is an equivalence relation on smooth curves through p : it does not depend on which chart is chosen and it is invariant under diffeomorphisms of M .

Another, equivalent way to define tangent vectors is as *derivations*. If X is the tangent vector corresponding to the curve γ through p , and $f: M \rightarrow \mathbb{R}$ is a scalar function on M , then

$$Xf = \left. \frac{d}{dt} \right|_{t=0} f(\gamma(t)). \quad (2.3)$$

Example. Let $p \in O(n)$. $T_p O(n)$ is the space of $n \times n$ matrices X such that $p^T X$ is skew-symmetric.

We also explain what it means for a tangent vector to vary smoothly. Let $\phi: U \rightarrow V$ be a chart, and let¹ $X: (p: M) \rightarrow T_p M$ be a map from a smooth manifold M to the tangent space at the point where it is evaluated. Then, for each $x \in V$, the tangent vector $X(\phi^{-1}(x)) \in T_{\phi^{-1}(x)} M$ corresponds to a velocity vector $\gamma'_{\phi^{-1}(x)} \in \mathbb{R}^n$. If this velocity vector is a smooth function of x , then X is called a *smooth vector field*. The vector space of smooth vector fields on a smooth manifold M is denoted $\mathcal{X}(M)$.

¹In this notation, we highlight that the target set depends on the input value. It should be read as “ $X: M \rightarrow \bigcup_{p \in M} T_p M$ such that $X(p) \in T_p M$ for every p ”.

2.1.1 Derivatives

We are now ready to start differentiating!

Definition 2.3. Let $f: M \rightarrow N$ be a smooth map between smooth manifolds. The *derivative*, or *differential*, $df_p: T_pM \rightarrow T_{f(p)}N$ at a point $p \in M$ is the map that takes a tangent vector $\gamma'(0)$ at p to the tangent vector $(f \circ \gamma)'(0)$ at $f(p)$.

On the level of smooth curves through p , df_p is just the map $\gamma \mapsto f \circ \gamma$. One verifies that definition 2.3 respects the equivalence relation by looking at γ and $f \circ \gamma$ in their respective charts. With this perspective, the chain rule $d(f \circ g)_p = df_{g(p)} \circ dg_p$ is almost a tautology.

2.1.2 Smooth submersions, immersions, and embeddings

It turns out to be very useful to study maps of *constant rank*, that is smooth maps $f: M \rightarrow N$ between smooth manifolds whose derivative have the same rank at every point.

Definition 2.4. A *smooth submersion* is a smooth map between smooth manifolds whose derivative is surjective everywhere.

Example. Consider the map that takes an orthogonal matrix $p \in O(n)$ and returns its first k columns. This is a smooth submersion $f: O(n) \rightarrow \text{St}(n, k)$, where $\text{St}(n, k)$ is the *Stiefel manifold*, since any smooth curve λ in $\text{St}(n, k)$ has a preimage $[\lambda \quad \lambda^\perp]$ in $O(n)$, where λ^\perp is an orthogonal basis completion. If one starts with an arbitrary basis completion to $\lambda(0)$ (which is always valid in some neighbourhood) and considers the steps taken in the Gram–Schmidt process, one verifies that such a basis completion can indeed be made smoothly.

Any map $f: M \rightarrow N$ defines an equivalence relation on its domain via $f(p) = f(q) \implies p \sim q$, but the equivalence relation induced by a surjective smooth submersion is also compatible with the smooth structure. In fact, this is a characteristic property.

Proposition 2.5 (Characteristic property for surjective smooth submersions). *Let $f: M \rightarrow N$ be a surjective smooth submersion. Then a map $F: N \rightarrow \mathbb{R}$ is smooth if and only if $F \circ f$ is smooth.*

Example. Consider again the map $f: O(n) \rightarrow \text{St}(n, k)$ from the previous example. The smooth structure of $\text{St}(n, k)$ is usually induced by embedding it as a smooth submanifold of $\mathbb{R}^{n \times k}$, but it can just as well be induced by demanding

that f is a submersion. This leads to the quotient representation $\text{St}(n, k) = \text{O}(n)/\text{O}(n - k)$. The two smooth structures on $\text{St}(n, k)$ are diffeomorphic.

Definition 2.6. A *smooth immersion* is a smooth map between smooth manifolds whose derivative is injective everywhere.

Example. The product

$$\begin{aligned} f: \mathbb{R}_{>0} \times S^{m-1} \times S^{n-1} &\rightarrow \mathbb{R}^{m \times n}, \\ (\lambda, x, y) &\rightarrow \lambda xy^\top \end{aligned} \tag{2.4}$$

is a smooth immersion. Its image is the set of rank-1 matrices and its derivative is

$$\begin{aligned} df_{(\lambda, x, y)}: \mathbb{R} \times T_x S^{m-1} \times T_y S^{n-1} &\rightarrow \mathbb{R}^{m \times n} \\ (\nu, u, v) &\mapsto \nu xy^\top + uy^\top + xv^\top. \end{aligned} \tag{2.5}$$

Here, to make sense of the products, we think of S^{m-1} and S^{n-1} as subsets of \mathbb{R}^m and \mathbb{R}^n respectively, and of $T_x S^{m-1}$ and $T_y S^{n-1}$ as linear subspaces of \mathbb{R}^m and \mathbb{R}^n respectively.

Definition 2.7. A *smooth embedding* is a smooth immersion that is also a homeomorphism onto its image.

The obstruction that keeps f from being an embedding in the previous example is injectivity: $f(\lambda, x, y) = f(\lambda, -x, -y)$. We say that $\mathbb{R}_{>0} \times S^{m-1} \times S^{n-1} \rightarrow \mathbb{R}^{m \times n}$ is a *double covering* of the set of rank-1 matrices. To get around this, we can impose an equivalence relation on f 's domain: $(\lambda, x, y) \sim (\lambda', x', y')$ if $\lambda = \lambda'$ and $(x, y) = \pm(x', y')$.

$$f: \frac{\mathbb{R}_{>0} \times S^{m-1} \times S^{n-1}}{\sim} \rightarrow \mathbb{R}^{m \times n} \tag{2.6}$$

is a smooth embedding of rank-1 matrices into Euclidean space. Let \mathcal{M}_1 denote f 's image. Rather than think of the smooth structure on \mathcal{M}_1 in terms of $(\mathbb{R}_{>0} \times S^{m-1} \times S^{n-1})/\sim$, we can just think of \mathcal{M}_1 as a submanifold of $\mathbb{R}^{m \times n}$. In general, we say that $M \subset N$ is an *embedded submanifold*, or just *submanifold*, if the inclusion map $\iota: M \rightarrow N$ is a smooth embedding.

2.2 Riemannian manifolds

The main references for this section are Lee's [Lee18] book on Riemannian manifolds and Cheeger and Ebin's [CE08] book on comparison theory.

Informally, a Riemannian manifold is smooth manifold that is also a metric space, along with some compatibility conditions between the smooth structure and the distance function. The way to make this rigorous is via an inner product on the tangent space.

Definition 2.8. A *Riemannian manifold* M is a smooth manifold, together with an inner product $\langle \cdot, \cdot \rangle_p$ on each tangent space $T_p M$ such that, for any two smooth vector fields X and Y ,

$$\begin{aligned} f: M &\rightarrow \mathbb{R}, \\ p &\mapsto \langle X(p), Y(p) \rangle_p \end{aligned} \tag{2.7}$$

is smooth.

We say that M is equipped with a *Riemannian structure*. From here, we can define a special class of curves, called *geodesics*, in two different ways.

First, note that the inner product induces a notion of lengths of curves on M :

$$\ell(\gamma) = \int_{\gamma} \sqrt{\langle \gamma'(t), \gamma'(t) \rangle_{\gamma(t)}} dt. \tag{2.8}$$

Importantly, ℓ is invariant under reparametrizations, unlike if we for example integrate $\langle \gamma', \gamma' \rangle$ without the square root. We can define a distance function on M : if $p, q \in M$, then

$$d(p, q) = \min_{\gamma} \ell(\gamma), \tag{2.9}$$

where γ ranges over the set of smooth curves from p to q .

Remark. This set might be empty, or otherwise not have a minimum.

Analogously to how one might solve for the minimum of a univariate real function f by setting $f'(t) = 0$, we can solve the variational problem (2.9) via the *Euler–Lagrange equation* $\delta\ell = 0$. See Gelfand and Silverman [GS+00] for a precise definition of the functional derivative. The solutions to the Euler–Lagrange equation are the critical points of ℓ . They are never local maxima, because we can always make a curve longer by adding loops. Those critical points are called *geodesics*.

Second, we define yet another type of derivative on smooth manifolds.

Definition 2.9. An *affine connection* on M is a map

$$\nabla: \mathcal{X}(M) \times \mathcal{X}(M) \rightarrow \mathcal{X}(M) \tag{2.10}$$

satisfying

1. $C^\infty(M)$ -linearity in the first argument: $\nabla_{fX+gY}Z = f\nabla_XZ + g\nabla_YZ$,
2. \mathbb{R} -linearity in the second argument: $\nabla_X(aY + bZ) = a\nabla_XY + b\nabla_XZ$,
3. the product rule: $\nabla_X(fY) = f\nabla_XY + (Xf)Y$,

for all $a, b \in \mathbb{R}$; $f, g \in C^\infty(M)$; and $X, Y, Z \in \mathcal{X}(M)$.

An affine connection defines a notion of *parallel transport* of X along a curve γ via the ODE

$$\nabla_{\gamma'(t)}X = 0. \quad (2.11)$$

If $\gamma'(t)$ satisfies its own parallel transport equation,

$$\nabla_{\gamma'(t)}\gamma'(t) = 0, \quad (2.12)$$

then γ is called a *geodesic*.

In order for these two notions of geodesic to coincide, we need a compatibility condition between the connection and the Riemannian metric.

Definition 2.10. The *Levi-Civita connection* is the unique affine connection satisfying

1. the product rule: $X\langle Y, Z \rangle = \langle \nabla_XY, Z \rangle + \langle Y, \nabla_XZ \rangle$,
2. the torsion condition: $\nabla_XY - \nabla_YX = [X, Y]$.

The combination of these two perspectives has very useful applications, especially in Riemannian comparison theory. For a proof that these two perspectives are equivalent, see [Lee18, corollary 6.7].

Geodesics provide a special type of chart, called *normal coordinates*. If M is a Riemannian manifold and let $p \in M$, then there are neighbourhoods $S \subset T_pM$ of 0 and $S' \subset M$ of p such that the following map is well-defined and invertible,

$$X \mapsto \gamma(1), \quad (2.13)$$

where γ is the constant-speed geodesic satisfying

$$\gamma(0) = p, \quad (2.14)$$

$$\gamma'(0) = X. \quad (2.15)$$

This map is denoted \exp_p and is called *exponential map*, or *manifold exponential*. Similarly, its inverse is denoted \log_p and is called *logarithmic map*, or *manifold logarithm*. The construction of the exponential map via solutions to an initial value problem (IVP) guarantees that it is a diffeomorphism in a neighbourhood of 0, and thus is a chart [Lee18, lemma 5.10].

2.2.1 Curvature

If M is a two-dimensional Riemannian manifold, then we can define the *Gaussian curvature* K at a point $p \in M$ in several ways. One way is to define it via the distance function d ,

$$K = \lim_{r \rightarrow 0} 3 \frac{2\pi r - C(r)}{\pi r^3}, \quad (2.16)$$

where $C(r)$ is the length of the “circle” $\{q \in M \text{ s.th. } d(p, q) = r\}$. Another option is to define K from the affine connection,

$$K = \langle \nabla_X \nabla_Y Y - \nabla_Y \nabla_X Y - \nabla_{[X, Y]} Y, X \rangle, \quad (2.17)$$

where X, Y is an orthonormal basis (ONB) for $T_p M$. It is possible to show that (2.17) does not depend on which ONB is chosen, and further does not depend on which smooth extension of X and Y is used.

If M is not two-dimensional, consider a two-dimensional linear subspace Π of $T_p M$. Then $\Sigma = \exp_p(\Pi)$ is a two-dimensional submanifold of M . The Gaussian curvature of Σ at p is called the *sectional curvature* of M along Π .

A complete and simply connected manifold with constant sectional curvature H is called a *model manifold*. It is uniquely determined by H and its dimension.

2.2.2 Riemannian submersions, immersions, and embeddings

Suppose $\pi: M \rightarrow N$ is a smooth submersion between Riemannian manifolds and that $\pi(p) = q$. Then $\pi^{-1}(q)$ is a smooth submanifold of M by the implicit function theorem [Lee13, theorem C.40]. The tangent space $T_p M$ can be split into two parts: the part tangent to $\pi^{-1}(q)$, called *vertical space* and denoted V_p , and the part orthogonal to the vertical space, called *horizontal space* and denoted H_p . We say that π is a *Riemannian submersion* if $d\pi$ preserves the inner product on H_p for all p .

Similarly to how a smooth submersion is a compatibility condition between the smooth manifolds, we can induce a Riemannian structure on M from a

Riemannian structure on N by demanding that some smooth submersion from $M \rightarrow N$ is also a Riemannian submersion.

A *Riemannian immersion* is a map between Riemannian manifolds whose derivative is everywhere injective (smooth immersion) and preserves the inner product. A *Riemannian embedding* is a Riemannian immersion that is also a homeomorphism to its image.

2.2.3 Toponogov's theorem

Riemannian comparison theory studies inequalities involving geometric quantities. A good reference is the book by Cheeger and Ebin [CE08], or the chapter by Karcher [Kar87]. Many comparison theorems relate properties of geodesics to curvature bounds. *Toponogov's theorem* is such a theorem, bounding side lengths of geodesic triangles defined on manifolds with lower bounded sectional curvature. It is illustrated in figure 2.1. We say that a subset of a Riemannian manifold is *geodesically convex* if, for any two points in the subset that are connected by a geodesic on the manifold, that geodesic lies in the subset as well.




Figure 2.1: Proposition 2.11 gives conditions for which the length of the opposite (dashed) side in M is bounded by the length of the opposite (dashed) side in N . Equal-length geodesics and equal angles have the same number of notches.

Proposition 2.11 (Toponogov's theorem [Pet16, Theorem 12.2.2 hinge version]). *Let M be a geodesically convex Riemannian manifold with sectional curvature bounded from below by some constant H . Let $\gamma_1, \gamma_2: [0, 1] \rightarrow M$ be minimal geodesics such that $\gamma_1(0) = \gamma_2(0)$. If $H > 0$, also assume that $d_M(\gamma_1(0), \gamma_1(1)), d_M(\gamma_2(0), \gamma_2(1)) \leq \frac{\pi}{\sqrt{H}}$. Consider a model manifold N of constant sectional*

curvature H and let $\lambda_1, \lambda_2: [0, 1] \rightarrow N$ be geodesics such that

$$\lambda_1(0) = \lambda_2(0), \quad (2.18)$$

$$d_N(\lambda_1(0), \lambda_1(1)) = d_M(\gamma_1(0), \gamma_1(1)), \quad (2.19)$$

$$d_N(\lambda_2(0), \lambda_2(1)) = d_M(\gamma_2(0), \gamma_2(1)). \quad (2.20)$$

and the angle where they meet satisfies $\angle(\lambda_1, \lambda_2) = \angle(\gamma_1, \gamma_2)$. Then

$$d_M(\gamma_1(1), \gamma_2(1)) \leq d_N(\lambda_1(1), \lambda_2(1)). \quad (2.21)$$

We also mention a ‘‘companion theorem’’ to Toponogov’s, illustrated in figure 2.2.




Figure 2.2: Proposition 2.12 gives conditions for which the length of the sector (dashed) in M is bounded by the length of the sector (dashed) in N . Equal-length geodesics and equal angles have the same number of notches/lines.

Proposition 2.12 ([Kar87, Theorem 4.1]). *Let M be a geodesically convex Riemannian manifold with sectional curvature bounded from above by some constant C . Let $\gamma_1, \gamma_2, \lambda_1, \lambda_2$ satisfy the same conditions as in proposition 2.11, and let N be a model manifold of constant sectional curvature C . Let γ_3 be an angle sector of $\angle(\gamma_1, \gamma_2)$ and let λ_3 be a similar angle sector of $\angle(\lambda_1, \lambda_2)$. Then*

$$d_M(\gamma_3(0), \gamma_3(1)) \leq d_N(\lambda_3(0), \lambda_3(1)). \quad (2.22)$$

The advantage of looking at geodesic triangles in the model manifold is that there are spherical and hyperbolic versions of the trigonometric identities. These identities are used later to find explicit error bounds when approximating maps into manifolds.

Lemma 2.13 (Reid [RS05, Sections 3.2 and 3.10]). *Let N be a manifold of constant sectional curvature H and consider a geodesic triangle on N with side*

lengths A, B, C and opposite angles a, b, c respectively. Then

$$\cos(C\sqrt{H}) = \cos(A\sqrt{H})\cos(B\sqrt{H}) + \sin(A\sqrt{H})\sin(B\sqrt{H})\cos c \quad (2.23)$$

if $H > 0$, and

$$\cosh(C\sqrt{|H|}) = \cosh(A\sqrt{|H|})\cosh(B\sqrt{|H|}) - \sinh(A\sqrt{|H|})\sinh(B\sqrt{|H|})\cos c \quad (2.24)$$

if $H < 0$.

2.2.4 Klingenberg's theorem

Klingenberg's theorem is a statement about how far geodesics can be assumed to be minimizing. Informally, it states that geodesics can fail to be minimizing for two reasons: large positive curvature or topological obstructions.

Definition 2.14. Let p be a point on a Riemannian manifold M . The *injectivity radius* is the biggest number i_p such that the exponential map is invertible on the geodesic ball $B_i(p)$. The *global injectivity radius* (often just *injectivity radius*) is the infimum i over all i_p .

Theorem 2.15 (Klingenberg's theorem [Pet16, lemma 6.4.7]). *Let M be a compact Riemannian manifold with sectional curvature bounded from above by some constant C , and let $p \in M$. Then*

$$i_p \geq \min \left\{ \frac{\pi}{\sqrt{C}}, \frac{l_p}{2} \right\}, \quad \text{and} \quad i \geq \min \left\{ \frac{\pi}{\sqrt{C}}, \frac{l}{2} \right\}, \quad (2.25)$$

where l_p is the shortest geodesic loop through p , and l is the infimum over all l_p .

Remark. If C is negative then π/\sqrt{C} should be interpreted as ∞ .

There is also a formulation that does not assume that M is compact, but instead assumes the curvature is also lower bounded by some positive constant. See for example [CE08, corollary 5.7]. In such a situation we say that the curvature is *pinched*. There is a lot of comparison theory about pinched manifolds. Here, we just mention that it is a very restrictive condition. For example if $C/4 < K \leq C$, then M is diffeomorphic to a sphere.

2.3 Lie groups and homogeneous spaces

The main references for this sections are Lee's [Lee13] book on smooth manifolds and O'Neill's [ONe83] book on semi-Riemannian geometry. Arvanitogeor-

gos' [Arv03] book on Lie groups and homogeneous spaces is also a good and concise reference for the material.

Definition 2.16. A *Lie group* is a set G such that

- G is a smooth manifold.
- G is a group.
- The group operation $*$: $G \times G \rightarrow G$ and inversion \cdot^{-1} : $G \rightarrow G$ are smooth maps.

Lie groups are a standard tools in Riemannian geometry. By defining the metric via compatibility conditions with the group operation, we get a lot of examples of Riemannian manifolds with nice properties.

A *Lie algebra* is a vector space \mathfrak{g} together with an antisymmetric bilinear operation called a *Lie bracket* $[\cdot, \cdot]: \mathfrak{g} \times \mathfrak{g} \rightarrow \mathfrak{g}$, which satisfies a kind of product rule called the *Jacobi identity*: $[v, [u, w]] = [[v, u], w] + [u, [v, w]]$. One can check that the matrix commutator is a Lie bracket, hence the overlapping notation.

The space $\mathcal{X}(M)$ of vector fields on a smooth manifold M forms a Lie algebra with $[X, Y]f = X(Yf) - Y(Xf)$. Moreover, if $M = G$ is a Lie group then we can restrict to so-called *left-invariant* vector fields, i.e. vector fields that are invariant under left translation. This space together with the commutator then forms a Lie algebra² with the same dimension as G . We call this G 's Lie algebra.

A fundamental result in Lie group theory is *Ado's theorem*, which states that every finite-dimensional Lie algebra can be realized as a matrix algebra, where the Lie bracket is the matrix commutator. Some examples are

- the space $\mathfrak{gl}(n)$ of $n \times n$ matrices,
- the space $\mathfrak{o}(n)$ of skew-symmetric $n \times n$ matrices,
- the space of upper-triangular $n \times n$ matrices.

Remarkably, Lie groups are in large part determined by their Lie algebras. If \mathfrak{g} is a $n \times n$ matrix Lie algebra, then the matrix exponential maps from $\mathfrak{g} \rightarrow \text{GL}(n)$. Its image $G' = \exp(\mathfrak{g})$ is a Lie subgroup of $\text{GL}(n)$. If \mathfrak{g} is the Lie

²One readily checks that the commutator of two left-invariant vector fields is again left-invariant.

algebra of some simply connected Lie group G , then $G = G'$. More generally, G' is the universal covering group of G , i.e. the smallest simply connected covering group.

For our intents and purposes, Lie groups are synonymous with matrix groups. Some examples are

- the space $GL(n)$ of invertible $n \times n$ matrices,
- the space $O(n)$ of orthogonal $n \times n$ matrices,
- the space of invertible upper-triangular $n \times n$ matrices.

Note that these are not connected, and so are not simply connected, but that the identity component is simply connected. Examples of non-simply connected Lie groups are somewhat more esoteric. The simply connected group $SL(2)$ of 2×2 determinant 1 matrices is a double cover of the non-simply connected projective group $PSL(2) = SL(2)/\{1, -1\}$.

2.3.1 Smooth quotients

One major reason to why Lie groups are interesting is because of their actions on smooth manifolds. A *Lie group action* θ of a Lie group G on a smooth manifold M is a map that takes elements in G to diffeomorphisms of M , so that group multiplication in G corresponds to composition of diffeomorphisms on M . More formally, $\theta: G \rightarrow \text{Diff}(M)$ is equivariant, where $\text{Diff}(M)$ denotes the group of diffeomorphisms on M . We often write $g \cdot p$ for “the action of $g \in G$ on $p \in M$ ”, and $G \curvearrowright M$ for “ G has an action on M ”.

Theorem 2.17. *If $H \subset G$ is a closed Lie subgroup, then there is a unique smooth structure on G/H such that $\pi: G \rightarrow G/H$ is a smooth submersion.*

Remarkably, theorem 2.17 has a converse. An action $G \curvearrowright M$ is *transitive* if $M = G \cdot p$ for some $p \in M$. By G 's group structure, if that holds for any p it holds for all p . M is then called *homogeneous*.

Theorem 2.18. *If $G \curvearrowright M$ is a transitive smooth action then $M = G/H$ as a smooth manifold, where H is the subgroup of G that fixes some given $p \in M$.*

Remark. The subgroup H in theorem 2.18 is called *stabilizer*, or *isotropy subgroup*.

Theorems 2.17 and 2.18 equate homogeneous manifolds with Lie group quotients whose denominator is closed. The terms “homogeneous manifold” and “Lie group quotient” are thus used interchangeably.

This gives us a plethora of new examples of smooth manifolds.

Example. Consider the action of $O(n)$ on the set of k -dimensional linear subspaces of \mathbb{R}^n . The action is clearly transitive. Moreover, the stabilizer is the subgroup that fixes $\text{span}\{e_1, \dots, e_k\}$ and $\text{span}\{e_{k+1}, \dots, e_n\}$: $O(k) \times O(n-k)$. This is the *Grassmann manifold*, $\text{Gr}(n, k) = O(n)/(O(k) \times O(n-k))$.

Example. Consider the action of $O(n)$ on the chain $V_1 \subset \dots \subset V_l$ of linear subspaces of \mathbb{R}^n with dimensions $k_1 \leq \dots \leq k_l$. Again, the action is clearly transitive. For simplicity, assume $k_1 = 0$ and $k_l = n$, then the stabilizer is $O(k_1) \times \dots \times O(k_l)$. This is the *flag manifold*, $F(k_1, \dots, k_l) = O(n)/(O(k_1) \times \dots \times O(k_l))$.

The flag manifold can be thought of as retaining more information than the Grassmann, but less than the Stiefel. See Ye, Wong, and Lim's [YWL22] recent paper, and the references therein, for an overview of the applications of flags.

It is also natural to consider similar quotients of non-orthogonal groups. Bendokat and Zimmermann [BZ21] for example study the *symplectic Grassmann* $\text{SpGr}(2n, 2k) = \text{Sp}(2n)/(\text{Sp}(2k) \times \text{Sp}(2(n-k)))$ and *symplectic Stiefel* $\text{SpSt}(2n, 2k) = \text{Sp}(2n)/\text{Sp}(2(n-k))$, consisting of $2k$ -dimensional symplectic subspaces and symplectic $2k$ -frames respectively. Similarly, one can similarly define the *symplectic flag manifold*.

2.3.2 Riemannian quotients

There is a natural Riemannian structure on smooth homogeneous spaces. It is defined via the Lie group quotient and a compatibility condition with the projection map. Such manifolds have nice properties, like their geometry looking roughly the same around every point, and later in section 5.4 they form a main class of examples where we are able to apply our approximation theorem.

Proposition 2.19. *Let G be a Lie group and H a Lie subgroup and assume G has a right-invariant Riemannian metric. Then there is a unique Riemannian metric on G/H such that*

$$\begin{aligned} \pi: G &\rightarrow G/H, \\ g &\mapsto gH \end{aligned} \tag{2.26}$$

is a Riemannian submersion.

Proposition 2.19 is obvious to experts in the subject. However, we were unable to find a reference for the precise statement. We therefore provide a short

proof in section 2.6 The idea is to identify the horizontal space with the tangent space of G/H . Note that a right-invariant metric on G does not necessarily induce a right-invariant metric on G/H . We hence make the following definition.

Definition 2.20. Let G be a Lie group with Lie algebra \mathfrak{g} , and let H be a Lie subgroup with Lie subalgebra \mathfrak{h} . G/H is *reductive* if there is an $\text{Ad}(H)$ -invariant³ linear subspace $\mathfrak{m} \subset \mathfrak{g}$, which is linearly independent of \mathfrak{h} .

Proposition 2.21 ([ONe83, Proposition 11.22]). *Let $M = G/H$ be reductive with $\text{Ad}(H)$ -invariant subspace \mathfrak{m} . Let $o = 1H$ be the coset of the identity. The condition that $d\pi : \mathfrak{m} \rightarrow T_o(M)$ is an isometry identifies $\text{Ad}(H)$ -invariant inner products on \mathfrak{m} with invariant metrics on M .*

We are thus interested in classifying which Lie group quotients are reductive. Whenever H is compact, G/H is reductive. For example, the space of full-rank symmetric matrices is the quotient $\text{GL}(n)/\text{O}(n)$ and is hence reductive. The set of rank- r symmetric matrices is the quotient

$$\text{GL}(n)/\left[\begin{array}{c} \text{GL}(n-r) \quad \mathbb{R}^{(n-r) \times r} \\ \text{GL}(r) \end{array}\right]. \quad (2.27)$$

The idea is that the denominator “quotients out the last $n-r$ columns” in a symmetric rank decomposition. It is only reductive in the trivial cases $r=0$ and $r=n$. The quotient $\text{GL}(n)/\text{GL}(r)$ is however always reductive.

Geodesics in G whose initial velocity is horizontal will keep having horizontal velocity, allowing us to define *horizontal geodesics*. From this also follows a one-to-one correspondence between horizontal geodesics in G through g and geodesics in M through $\pi(g)$ [CE08, proposition 3.31]. In terms of the manifold exponential, we can write

$$\exp_{\pi(g)}(d\pi_g X) = \pi(\exp_g(X)) \quad (2.28)$$

when $X \in T_g G$ is horizontal. This is useful because it allows us to lift geodesics in M as geodesics in G .

2.4 Algebraic geometry

Much of the material in this section can be found in Smith, Kahanpaa, Kekalainen, and Traves’ [Smi+04] introduction to algebraic geometry.

³Here, $\text{Ad}(g)$ is *adjoint action* $v \mapsto gvg^{-1}$ on \mathfrak{g} .

The central objects of study in algebraic geometry are *algebraic sets*, or *varieties*. The sphere, the unitary group, and rank-at-most- r matrices are examples. They are subsets of some \mathbb{C}^n that are also solution sets to polynomial equations. Much of the mathematics in algebraic geometry comes from a bridge between geometry and commutative algebra: on one hand varieties are sets of points in a vector space, while on the other hand the polynomials that define them form ideals in a commutative ring of polynomials. This bridge yields a rich theory of such sets.

The *Zariski topology* is a special topology on algebraic varieties, where closed sets are algebraic subvarieties. It is very different from the standard topology induced by open intervals on $\mathbb{C}^n = \mathbb{R}^{2n}$. Sets that are closed/open in the Zariski topology are closed/open in the standard topology, but the converse is not true in general. Notably, if $X \subset Y$ is a Zariski closed subset, then either $X = Y$ or X is a (Lebesgue) measure zero set. We thus say that a property is *generic* if it holds on a nonempty Zariski open set.

For our purposes, we would also like to relate the notion of variety to that of manifolds. Not all varieties are manifolds and not all embedded manifolds are varieties. However, some are both. Spheres and unitary groups are both varieties and manifolds. Rank at-most- r matrices have an open and dense (in the standard topology) subset which is a manifold: matrices with rank exactly r . To see that they are a variety, consider first that the set of matrices with rank at most $n - 1$ is the solution to $\det A = 0$. Then, it is straightforward to see that the set of matrices with rank at most r is set of matrices whose $(r - 1) \times (r - 1)$ minors vanish.

Section 2.5 is about tensors and tensor ranks. There, we will see that upper bounding the rank is similarly an algebraic condition for tensors. Because of this much of what we know about rank-restricted tensors comes from algebraic geometry.

We mention also that restricting from \mathbb{C}^n to \mathbb{R}^n is not an algebraic condition, and that the theory of real algebraic geometry is very different. In essence, this is because the reals are not algebraically closed. Therefore, we need to be careful when applying algebraic geometrical results in the real setting.

2.5 Tensors

The main references for this section are Landsberg's [Lan12] book on the algebraic geometry of tensors and Kolda and Bader's [KB09] review paper.

2.5.1 Characteristic property

A *tensor* is a multilinear map. Moreover, we can canonically identify multilinear maps with linear maps via the *tensor product*, \otimes , which we define with the following characteristic property.

Proposition 2.22 (Characteristic property for tensor products, [Gre78, section 1.4]). *Let A_1, \dots, A_d be real vector spaces. There is a real vector space $A_1 \otimes \dots \otimes A_d$ and a multilinear map $\phi: A_1 \times \dots \times A_d \rightarrow A_1 \otimes \dots \otimes A_d$ with the property that for each multilinear map $f: A_1 \times \dots \times A_d \rightarrow \mathbb{R}$ there exists a unique linear map $F: A_1 \otimes \dots \otimes A_d \rightarrow \mathbb{R}$ such that*

$$\begin{array}{ccc}
 A_1 \times \dots \times A_d & & \\
 \phi \downarrow & \searrow f & \\
 A_1 \otimes \dots \otimes A_d & \xrightarrow{F} & \mathbb{R}
 \end{array} \tag{2.29}$$

is a commutative diagram.

It is characteristic in the sense that $A_1 \otimes \dots \otimes A_d$ and ϕ are essentially unique. We could thus equivalently have defined a tensor as an element of a tensor product space. The factors A_i are called *modes*, and the number of modes, d , is called the *order* of the tensor.

From now, we restrict our attention to finite-dimensional vector spaces. Then, the *Kronecker product*, also denoted \otimes , is a realization of ϕ . If $a = (a_1, \dots, a_m) \in \mathbb{R}^m$ and $b = (b_1, \dots, b_n) \in \mathbb{R}^n$, then in block matrix notation, we have

$$a \otimes b = \begin{bmatrix} a_1 b \\ \vdots \\ a_m b \end{bmatrix} \in \mathbb{R}^{mn}. \tag{2.30}$$

Furthermore, if $L: \mathbb{R}^m \rightarrow \mathbb{R}^n$ and $M: \mathbb{R}^p \rightarrow \mathbb{R}^q$ are linear operators, then $L \otimes M: \mathbb{R}^m \otimes \mathbb{R}^p \rightarrow \mathbb{R}^n \otimes \mathbb{R}^q$ is the linear operator that maps $a \otimes b$ to $L(a) \otimes M(b)$. In block matrix notation, we have

$$L \otimes M = \begin{bmatrix} L_{11}M & \dots & L_{1m}M \\ \vdots & & \vdots \\ L_{n1}M & \dots & L_{nm}M \end{bmatrix} \in \mathbb{R}^{nq \times mp}. \tag{2.31}$$

Greub [Gre78] is a good introduction to tensor products via the characteristic property, as well as to the *symmetric product* \odot and the *skew-symmetric*

product \wedge . Landsberg [Lan12] is a good introduction to the algebraic geometric perspective on tensors.

2.5.2 Penrose diagrams

A mathematical expression is a directed acyclic graph (DAG), where the nodes are functions, the incoming edges are input arguments, and the outgoing edges are output values. For example, the expression $e = \exp(x^2 + y + z)$ is shown in figure 2.3. Since function composition is associative, such a graph is unambiguous. Moreover, for a function that takes several arguments, it does not matter in which order the arguments are supplied. We thus define *partial application*: $f(x, y) = f(x, \cdot)(y) = f(\cdot, y)(x)$. These equalities are reflected by their graphs being the same.




Figure 2.3: A mathematical expression.

We now specialize to tensors. For these, partial application has a special interpretation. Consider for example linear maps $L: A \rightarrow B$ and $M: B \rightarrow C$. Equivalently,

$$L: B^* \rightarrow A^*, \quad (2.32)$$

$$L: A \otimes B^* \rightarrow \mathbb{R}, \quad (2.33)$$

$$L: A^* \otimes B, \quad (2.34)$$

and so on, where $A^* = \{f: A \rightarrow \mathbb{R} \text{ s.th. } f \text{ is linear}\}$ is the *dual* of A . The composition $M \circ L: A \rightarrow C$ is the partial application of M to the output of L . It is shown in figure 2.4. Equivalently, we can write, $M \circ L: A^* \otimes C$. More generally, if we are given two tensors, one in $A_1 \otimes \dots \otimes A_d \otimes B$ and one in $B^* \otimes C_1 \otimes \dots \otimes C_{d'}$, there is a natural way to combine them to a tensor in $A_1 \otimes \dots \otimes A_d \otimes C_1 \otimes \dots \otimes C_{d'}$ by setting $A = A_1 \otimes \dots \otimes A_d$ and $C = C_1 \otimes \dots \otimes C_{d'}$ in the previous sentence.

Now, fix some basis for A , B , and C . If L has components L_{ij} and M components M_{jk} , then $M \circ L$ has components $\sum_j L_{ij} M_{jk}$. Such a sum is called a *contraction*. To summarize, a (partial) function application corresponds to an inner edge in the graph, which in turn corresponds to a contraction in a basis.



Figure 2.4: Composition of linear maps.

Lastly, we further specialize to tensors on inner product spaces. While finite-dimensional linear spaces are always isomorphic to their duals, that isomorphism is not a priori canonical. An inner product singles out one such isomorphism. For our graphs, this implies that the direction of edges is no longer meaningful. We are thus finally able to define *Penrose diagram*. It is an undirected graph where each node is a tensor and each edge is an inner product space.

We will later see some examples of Penrose diagrams, namely figures 2.5 to 2.7.

2.5.3 CP decomposition

A matrix that is the outer product of two vectors has rank 1, unless it is zero. Similarly, we say that a tensor that is an image of a tensor product has *rank 1*, unless it is zero. The set of rank-1 tensors, together with the zero tensor is an algebraic variety, called the *Segre variety*. The *Canonical Polyadic (CP) rank*, or *tensor rank*, of a tensor $T \in \mathbb{R}^{n_1} \otimes \cdots \otimes \mathbb{R}^{n_d}$ is the smallest r such that T is the sum of r rank-1 tensors:

$$T = \sum_{j=1}^r v_1^j \otimes \cdots \otimes v_d^j. \quad (2.35)$$

Such a (minimal) expression is called a *CP decomposition*, or *tensor rank decomposition*. If we arrange the factors in matrices $V_i = [v_i^1 \ \cdots \ v_i^r]$, then the notation

$$T = \llbracket V_1, \dots, V_d \rrbracket \quad (2.36)$$

can also be used. We can think of (2.35) as the Penrose diagram in figure 2.5. For matrices, their CP rank is just their matrix rank.

A CP decomposition is *essentially unique* if it is unique up to reordering the terms and scaling the factors with numbers that multiply to 1. For matrices, their rank decomposition is never essentially unique since $M = LR = LA^{-1}AR$ for




Figure 2.5: Penrose diagram for the CP decomposition. Edges are labeled with the vector space they represent, and $V_i = [v_i^1 \ \dots \ v_i^r]$. I is the diagonal tensor $I_{i_1 \dots i_d} = 1$ if $i_1 = \dots = i_d$ else $I_{i_1 \dots i_d} = 0$.

any $A \in \text{GL}(r)$. However, the following result states that the CP decomposition of higher order tensors can be unique under certain conditions.

Theorem 2.23 (Kruskal's theorem). *Let $T \in \mathbb{R}^{n_1} \otimes \dots \otimes \mathbb{R}^{n_d}$ be CP rank r with CP decomposition (2.35). Define k_i , for each mode i , so that every size k_i subset of the factors v_i^1, \dots, v_i^r is linearly independent. If*

$$2r + d - 1 \leq \sum_{i=1}^d k_i, \quad (2.37)$$

then T 's CP decomposition is essentially unique.

Unlike the matrix case, upper bounding the CP rank is not an algebraic condition for tensors. The reason is that such a bound is not a closed condition. Consider the sequence

$$T_N = N \left(a + \frac{1}{N} b \right) \otimes \left(a + \frac{1}{N} b \right) \otimes \left(a + \frac{1}{N} b \right) - N a \otimes a \otimes a \quad (2.38)$$

$$\begin{aligned} &= a \otimes a \otimes b + a \otimes b \otimes a + b \otimes a \otimes a \\ &\quad + \frac{1}{N} (a \otimes b \otimes b + b \otimes a \otimes b + b \otimes b \otimes a) + \frac{1}{N^2} b \otimes b \otimes b \end{aligned} \quad (2.39)$$

$$\xrightarrow{N \rightarrow \infty} a \otimes a \otimes b + a \otimes b \otimes a + b \otimes a \otimes a. \quad (2.40)$$

It is a sequence of rank-2 tensors whose limit is rank-3 [SL08, Theorem 1.1]. Tensors in the closure (standard topology) of CP rank- r tensors are said to have *border rank* r . This closure is an algebraic variety called a *secant variety* [Lan12, Chapter 4].

We stress how different tensor ranks are from their matrix analogues. Especially, many problems that are straightforward to solve for matrices are intractable for higher order tensors. We mention the paper “Most tensor problems are NP-hard” by Hillar and Lim [HL13]. For example, determining the exact rank of a tensor is NP-hard.

Consider moreover the *symmetric CP decomposition* of a symmetric tensor. It is the minimal number r of symmetric rank-1 terms that the tensor can be written as:

$$P = \sum_{j=1}^r v^j \otimes \cdots \otimes v^j. \quad (2.41)$$

It was conjectured by Comon, Golub, Lim, and Mourrain [Com+08] that the symmetric CP rank of a symmetric tensor is always the same as its CP rank. Currently, the conjecture is open⁴, although it has been proved for sufficiently small rank and some other special cases. See the survey by Casarotti, Massarenti, and Mella [CMM18] for an overview.

Symmetric tensors are especially interesting because they correspond to polynomials. A polynomial is *homogeneous*⁵ if each term has the same degree. In this case, we can identify the set of linear homogeneous polynomials in n variables with an n -dimensional vector space V by letting each independent variable be a basis vector. The space of degree- d homogeneous polynomials is then $V^{\odot d}$.

Similarly to the Segre variety, the set of rank-1 symmetric tensors together with the zero tensor is an algebraic variety, called the *Veronese variety*. More generally, the set of rank-1 tensors in $V_1^{\odot k_1} \otimes \cdots \otimes V^{\odot k_d}$ together with the zero tensor is called the (k_1, \dots, k_d) th *Segre–Veronese variety*.

As mentioned, the problem of finding a CP decomposition for an arbitrary tensor is NP-hard. See for example Håstad [Hås89] or Hillar and Lim [HL13]. One interpretation of this is that there is no efficient direct decomposition algorithm.

⁴Shitov [Shi18] published a counterexample to the conjecture, but an erratum was later added by Draisma [Dra24]. Exciting times!

⁵No relation to the *homogeneous* in “homogeneous manifold”.

The most common way to CP decompose a given tensor $T: \mathbb{R}^{n_1} \otimes \cdots \otimes \mathbb{R}^{n_d}$ in practice is to formulate it as an optimization problem. The natural objective function to consider is

$$f(v_1^1, \dots, v_d^r) = \left\| T - \sum_{j=1}^r v_1^j \otimes \cdots \otimes v_d^j \right\|_{\mathbb{F}}^2. \quad (2.42)$$

Alternating least squares (ALS) is a standard algorithm to solve this problem. It minimizes the objective by 1) fixing all variables except v_1^1, \dots, v_1^r and minimizes with respect to these (this is a linear least squares problem), 2) fixing all variables except v_2^1, \dots, v_2^r , and so on. Most implementations are based on the review paper by Kolda and Bader [KB09]. `TensorLy` is a Python tensor library with a machine learning focus. Its default CP decomposition algorithm implements ALS from Kolda and Bader. `TensorToolbox` [Bre+23] is a Matlab toolbox implementing a host of different structured tensor types, along with corresponding decomposition methods. Its standard method for CP decomposition is also ALS from Kolda and Bader. Similarly, `TensorToolbox.jl` is a Julia library that aims to port types and methods from Matlab's `TensorToolbox`. It also implements ALS from Kolda and Bader.

ALS methods are however no longer considered state of the art. Instead, so-called *all at once* methods that optimize all factors simultaneously have the best performance. We mention the gradient descent methods by Acar, Dunlavy, and Kolda [ADK11]; by Phan, Tichavský, and Cichocki [PTC13]; by Sorber, Van Barel, and De Lathauwer [SVD13]; and by Vermeylen, Vervliet, and De Lathauwer [VVD25]. We also mention the Riemannian trust region method by Breiding and Vannieuwenhoven [BV18], which treats the CP decomposition problem as an optimization problem on a product of *Segre manifolds*.

Note however that we need $\mathcal{O}(n_1 \cdots n_d)$ basic operations just to evaluate the objective function (2.42). One could say that the complexity is exponential in d . This explosion in the number of operations is known as *the curse of dimensionality*. A large part of current day applied research is about how to overcome it.

The following approach is one way to tackle the curse. We can make a simple modification to (2.42) so that we do not have to look at every entry in the tensor. Let I be a subset of T 's indices and consider the seminorm

$$\|A\|_{\mathbb{F}, I} = \sqrt{\sum_{(i_1, \dots, i_d) \in I} A_{i_1 \dots i_d}^2}. \quad (2.43)$$

That is, we are only summing over a subset of the entries. With such a seminorm, we can replace (2.42) with

$$f(v_1^1, \dots, v_d^r) = \left\| T - \sum_{j=1}^r v_1^j \otimes \dots \otimes v_d^j \right\|_{F,I}^2, \quad (2.44)$$

and control how many operations are needed with the size of I .

This approach can be especially profitable if not all entries of T are known. Therefore, these types of algorithms are often considered in the context of *tensor completion*. The more general problem of minimizing an arbitrary objective function to find an approximate tensor decomposition is sometimes called *generalized tensor decomposition*. These techniques are detailed in a paper by Kolda and Hong [KH20], and in a review by Hong, Kolda, and Duersch [HKD20].

While NP-hardness prohibits nice global convergence results for ALS-type algorithms, it is still possible to show *local convergence*. That is, if the initial guess is sufficiently close to a minimum, we can guarantee that we converge to towards that minimum. We mention Uschmajew [Usc12] and Hu, Iwen, Needell, and Wang [Hu+25] who show local convergence under some mild assumptions. We also mention that Hu and Li [HL18], Wang and Chu [WC14], and Wang, Chu, and Yu [WCY15] are able to show some global convergence result under more restrictive assumptions, like only considering rank-1 approximants.

One major application of CP decompositions is *blind source separation* (BSS). Informally, BSS aims to identify a set of source signals, with only information about their mixture. For example, the act of trying to listen to a conversation in a crowded room with a lot of other conversations is a BSS problem. One fruitful approach is to encode the mixed signal as a tensor, to decompose it, and to interpret each rank-1 term as a source. Therefore, a lot of work on practical decomposition algorithms has been done in this context. We mention the book by Comon and Jutten [CJ10] and the survey by Cichocki, Mandic, De Lathauwer, Zhou, Zhao, Caiafa, and Phan [Cic+15].

2.5.4 Tucker decomposition

Out of the three tensor decompositions that we consider, the *Tucker decomposition* is perhaps the most ubiquitous.

Definition 2.24. The *multilinear* (or *Tucker*) rank of a tensor $T \in \mathbb{R}^{n_1} \otimes \cdots \otimes \mathbb{R}^{n_d}$ is the tuple (t_1, \dots, t_d) with⁶

$$t_i = \text{rank}(T: \mathbb{R}^{n_i} \rightarrow \mathbb{R}^{n_1} \otimes \cdots \widehat{\mathbb{R}^{n_i}} \cdots \otimes \mathbb{R}^{n_d}). \quad (2.45)$$

Proposition 2.25. If $T \in \mathbb{R}^{n_1} \otimes \cdots \otimes \mathbb{R}^{n_d}$ has multilinear rank (t_1, \dots, t_d) then T can be written as

$$T_{k_1 \dots k_d} = \sum_{\alpha_1=1}^{t_1} \cdots \sum_{\alpha_d=1}^{t_d} C_{\alpha_1 \dots \alpha_d} (G_1)_{k_1}^{\alpha_1} \cdots (G_d)_{k_d}^{\alpha_d}, \quad (2.46)$$

where $C \in \mathbb{R}^{t_1} \otimes \cdots \otimes \mathbb{R}^{t_d}$ and $G_i \in \mathbb{R}^{n_i} \otimes (\mathbb{R}^{t_i})^*$. Such an expression is called a Tucker decomposition. It is illustrated in figure 2.6.

Proof. By (2.45), we can use the matrix rank decomposition in each mode to arrive at (2.46). \square

Also, conversely, a tensor admitting a Tucker decomposition with inner dimensions t_1, \dots, t_d has multilinear rank at most (t_1, \dots, t_d) . If the inner dimensions and multilinear rank are equal, then the G_i have maximal rank. Contrast this to TT tensors, where it is possible to have TT rank equal to the inner dimensions while some F_i has nonmaximal multilinear rank.

If the G_i have maximal rank, we could without loss of generality require that they be orthogonal. Many authors do this, but the stabilizer in chapter 3 will be easier to derive if we do not.

In proposition 2.25, we used the matrix rank decomposition recursively, and since that decomposition is unique up to a change of basis in the inner vector space, we immediately arrive at the following.

Proposition 2.26. The Tucker decomposition is unique up to a change of basis in $\mathbb{R}^{t_1}, \dots, \mathbb{R}^{t_d}$. More precisely,

$$\begin{aligned} \sum_{\alpha_1=1}^{t_1} \cdots \sum_{\alpha_d=1}^{t_d} C'_{\alpha_1 \dots \alpha_d} (G'_1)_{k_1}^{\alpha_1} \cdots (G'_d)_{k_d}^{\alpha_d} = \\ \sum_{\alpha_1=1}^{t_1} \cdots \sum_{\alpha_d=1}^{t_d} C_{\alpha_1 \dots \alpha_d} (G_1)_{k_1}^{\alpha_1} \cdots (G_d)_{k_d}^{\alpha_d} \end{aligned} \quad (2.47)$$

⁶Here, the hat notation means that factor is excluded.




Figure 2.6: Penrose diagram for the Tucker decomposition (2.46). Edges are labeled with the vector space they represent.

iff there are matrices $U_1 \in \text{GL}(t_1), \dots, U_d \in \text{GL}(t_d)$ such that

$$(G'_i)_{k_i}{}^{\alpha_i} = \sum_{\beta=1}^{t_i} (G_i)_{k_i}{}^{\beta} (U_i)_{\beta}{}^{\alpha_i}, \quad i = 1, \dots, d \quad (2.48)$$

$$(C')_{\alpha_1 \dots \alpha_d} = \sum_{\beta_1=1}^{t_1} \dots \sum_{\beta_d=1}^{t_d} C_{\beta_1 \dots \beta_d} (U_1^{-1})^{\beta_1}{}_{\alpha_1} \dots (U_d^{-1})^{\beta_d}{}_{\alpha_d}. \quad (2.49)$$

Similarly to CP and TT decompositions, there are several approaches to compute the Tucker decomposition. For the optimization approach, we mention De Lathauwer, De Moor, and Vandewalle [DDV00b]; Eldén and Savas [ES09]; Savas and Lim [SL10]; Ishteva, De Lathauwer, Absil, and Van Huffel [Ish+09]; Ishteva, Absil, Van Huffel, and De Lathauwer [Ish+11]; and Kressner, Steinlechner, and Vandereycken [KSV14].

There is an algorithm, called *HOSVD*, introduced in [DDV00a], to compute an orthogonal Tucker decomposition from successive SVDs applied to different matricizations of the tensor in question. Here we just mention that the essential difference between this decomposition and TT-SVD is the following: If $T: \mathbb{R}^{n_1} \otimes \dots \otimes \mathbb{R}^{n_d}$, then TT-SVD applies the SVD to $T: \mathbb{R}^{n_1} \otimes \dots \otimes \mathbb{R}^{n_{k-1}} \rightarrow \mathbb{R}^{n_k} \otimes \dots \otimes \mathbb{R}^{n_d}$, while HOSVD applies the SVD to $T: \mathbb{R}^{n_k} \rightarrow \mathbb{R}^{n_1} \otimes \dots \otimes \widehat{\mathbb{R}^{n_k}} \dots \otimes \mathbb{R}^{n_d}$. This difference is mirrored in the difference between definition 2.27 and definition 2.24. We also mention an efficient compression technique, *sequentially truncated HOSVD* (ST-HOSVD), introduced by Vannieuwenhoven, Vandebril, and Meerbergen [VVM12]. Moreover, like TT-cross, there is an analogue

of HOSVD that uses pseudoskeletons instead of SVDs, called *higher order interpolatory decomposition*, introduced by Saibaba [Sai16].

Several software libraries implement Tucker decompositions. `TensorLy` implements an optimization algorithm called *higher order orthogonal iteration* (HOOI). Both Matlab's `TensorToolbox` and Julia's `TensorToolbox.jl` implements HOSVD and ST-HOSVD.

2.5.5 Tensor trains

The *tensor train* (TT) decomposition is another possible generalization of matrix rank decomposition to tensors. It was also introduced independently in quantum mechanics under the name *matrix product state* (MPS).

Definition 2.27. The *TT rank* of a tensor $T \in \mathbb{R}^{n_1} \otimes \cdots \otimes \mathbb{R}^{n_d}$ is the tuple (s_1, \dots, s_{d-1}) with⁷

$$s_i = \text{rank}(T: \mathbb{R}^{n_1} \otimes \cdots \otimes \mathbb{R}^{n_i} \rightarrow \mathbb{R}^{n_{i+1}} \otimes \cdots \otimes \mathbb{R}^{n_d}). \quad (2.50)$$

Like matrix rank, upper bounding the TT rank is an algebraic condition, and thus defines an algebraic variety.

Proposition 2.28 (Oseledets [Ose11, Theorem 2.1]). *If $T \in \mathbb{R}^{n_1} \otimes \cdots \otimes \mathbb{R}^{n_d}$ has TT rank (s_1, \dots, s_{d-1}) then T can be written as*

$$T_{k_1 k_2 \dots k_d} = \sum_{\alpha_1=1}^{s_1} \cdots \sum_{\alpha_{d-1}=1}^{s_{d-1}} (F_1)_{k_1 \alpha_1} (F_2)^{\alpha_1 k_2 \alpha_2} \cdots (F_d)^{\alpha_{d-1} k_d}, \quad (2.51)$$

where $F_i \in (\mathbb{R}^{s_{i-1}})^* \otimes \mathbb{R}^{n_i} \otimes \mathbb{R}^{s_i}$. Such an expression is called a TT decomposition. It is illustrated in figure 2.7.




Figure 2.7: Penrose diagram for the TT decomposition (2.51). Edges are labeled with the vector space they represent.

Note also that a converse of proposition 2.28 is true, that a TT decomposition with inner dimensions s_1, \dots, s_{d-1} has TT rank at most (s_1, \dots, s_{d-1}) .

⁷Recall from section 2.5.2 the different interpretations of multilinear maps as linear maps.

Remark. There is a slight difference in the definition of TT and MPS: the latter allows the last factor F_d to be connected also to the first factor F_1 . This creates a loop in the Penrose diagram. MPS is thus more agnostic to the specific ordering between the indices. On the other hand, restricting the rank of an MPS is not always an algebraic condition.

Kruskal's theorem is a result about the essential uniqueness of CP decompositions. There is a similar uniqueness theorem about the TT decomposition. Compared to the CP case, the degree of uniqueness is weaker but it does not depend on the rank.

Proposition 2.29. *The TT decomposition is unique up to a change of basis in $\mathbb{R}^{s_1}, \dots, \mathbb{R}^{s_{d-1}}$. More precisely,*

$$\begin{aligned} \sum_{\alpha_1}^{s_1} \cdots \sum_{\alpha_{d-1}}^{s_{d-1}} (F'_1)_{k_1 \alpha_1} (F'_2)^{\alpha_1}_{k_2 \alpha_2} \cdots (F'_d)^{\alpha_{d-1}}_{k_d} = \\ \sum_{\alpha_1}^{s_1} \cdots \sum_{\alpha_{d-1}}^{s_{d-1}} (F_1)_{k_1 \alpha_1} (F_2)^{\alpha_1}_{k_2 \alpha_2} \cdots (F_d)^{\alpha_{d-1}}_{k_d} \end{aligned} \quad (2.52)$$

iff there are matrices $U_1 \in \text{GL}(s_1), \dots, U_{d-1} \in \text{GL}(s_{d-1})$ such that

$$(F'_1)_{k_i \alpha_i} = \sum_{\beta=1}^{s_1} (F_1)_{k_1 \beta} (U_1^{-1})^\beta_{\alpha_1}, \quad (2.53)$$

$$(F'_i)^{\alpha_{i-1}}_{k_i \alpha_i} = \sum_{\beta=1}^{s_{i-1}} \sum_{\gamma=1}^{s_i} (U_{i-1})^\beta_{\alpha_{i-1}} (F_i)^\beta_{k_i \gamma} (U_i^{-1})^\gamma_{\alpha_i}, \quad 2 \leq i \leq d-1, \quad (2.54)$$

$$(F'_d)^{\alpha_{d-1}}_{k_d} = \sum_{\beta=1}^{s_{d-1}} (U_{d-1})^\beta_{\alpha_{d-1}} (F_d)^\beta_{k_d}. \quad (2.55)$$

Uschmajew and Vandereycken [UV13, proposition 3] show that this holds for Hierarchical Tucker decompositions, of which the TT decomposition is a special case.

To compute a TT decomposition, one option is of course to define an objective function in the same way as (2.42) or (2.44). Holtz, Rohwedder, and Schneider develop an ALS algorithm for tensor trains [HRS12b], while Rohwedder and Uschmajew [RU13] and Oseledets, Rakhuba, and Uschmajew [ORU18; ORU23] derive local convergence results.

In contrast to CP decompositions, there is a host of direct methods for computing TT decompositions. We mention two of them here. The most straightforward method computes successive matrix rank decompositions of *unfolding matrices*. If $T \in \mathbb{R}^{n_1} \otimes \cdots \otimes \mathbb{R}^{n_d}$, then its k th unfolding is the corresponding linear map

$$T_{(k)}: \mathbb{R}^{n_1} \otimes \cdots \otimes \mathbb{R}^{n_k} \rightarrow \mathbb{R}^{n_{k+1}} \otimes \cdots \otimes \mathbb{R}^{n_d}. \quad (2.56)$$

Assume $\text{ttrank} T = (s_1, \dots, s_{d-1})$. By definition, for each k , there is a rank decomposition

$$T_{(k)} = LR, \quad (2.57)$$

where L is an $(n_k \cdots n_1) \times s_k$ matrix and R is an $s_k \times (n_{k+1} \cdots n_d)$ matrix. We can also interpret L and R as unfoldings of tensors, whose TT ranks satisfy

$$\text{ttrank}(L: \mathbb{R}^{n_1} \otimes \cdots \otimes \mathbb{R}^{n_k} \otimes \mathbb{R}^{s_k}) \leq (s_1, \dots, s_{k-1}), \quad (2.58)$$

$$\text{ttrank}(R: \mathbb{R}^{s_k} \otimes \mathbb{R}^{n_{k+1}} \otimes \cdots \otimes \mathbb{R}^{n_d}) \leq (s_k, \dots, s_{d-1}). \quad (2.59)$$

To understand why, consider first R , and note that L has a left inverse L' . Thus (2.57) becomes $R = L'T_{(k)}$ and the l th unfolding of R is

$$R_{(l)} = (L' \otimes I \otimes \cdots \otimes I)T_{(k+l)}. \quad (2.60)$$

From this follows (2.59) since the rank of a product of two matrices is bounded by the either of the ranks of the factors. Similarly for L shows (2.58). Using the rank decomposition recursively thus yields a TT decomposition, also showing that the inequalities are actually equalities. This method is called *TT-SVD* [Ose11]. Unfortunately, since it computes decompositions of unfoldings of the entire tensor, it suffers badly from the curse of dimensionality.

A more computationally feasible approach is based on *pseudoskeletons*, also called *adaptive cross approximation* (ACA) or *CUR decompositions*. To briefly explain the idea, let A be a rank- r $m \times n$ matrix and let \hat{A} be a nonsingular $r \times r$ submatrix. Let $I = (i_1, \dots, i_r)$ be \hat{A} 's row indices in A , and $J = (j_1, \dots, j_r)$ be its column indices and define⁸ $L = A_{:,I}$, $R = A_{J,:}$. Then

$$A = L\hat{A}^{-1}R. \quad (2.61)$$

⁸Here, we use matrix indices similarly to Matlab, so that $A_{:,I}$ is the matrix

$$\begin{bmatrix} A_{1i_1} & \cdots & A_{1i_r} \\ \vdots & & \vdots \\ A_{ni_1} & \cdots & A_{ni_r} \end{bmatrix}.$$

We have achieved a compact decomposition of the matrix A . Notably, we have done so without sampling all of its entries. The computational complexity $r \ll m, n$ is $\mathcal{O}(mr + nr + r^3)$. Compare this to an SVD, which is $\mathcal{O}((m+n)n^2)$.

The idea of pseudoskeletons is now that if A is not exactly rank r , but is well approximated by a rank- r matrix then (2.61) is an approximate equality. We do not go into details about the approximation properties of such a decomposition, but just mention that if \hat{A} is chosen to have maximal determinant, then the CUR approximation is quasi-optimal in the sense that

$$\left\| A - L\hat{A}R \right\|_{\max} \leq (r+1)\sigma_{r+1}, \quad (2.62)$$

where σ_{r+1} is the $(r+1)$ th singular value of A . The problem of maximizing \hat{A} 's determinant is called *submatrix selection*. For an explanation of why maximizing the determinant produces a quasi-optimal CUR, we refer to the classic paper by Goreinov, Tyrtyshnikov, and Zamarashkin [GTZ97], and to the book chapters by Goreinov, Oseledets, Savostyanov, Tyrtyshnikov, and Zamarashkin [Gor+10] and by Goreinov and Tyrtyshnikov [GT01]. The functional variant, *adaptive cross approximation*, is introduced by Bebendorf [Beb00]. We also mention the recent papers by Osinsky [Osi25], and by Cortinovis and Kressner [CK25], on a fast algorithm that is provably quasi-optimal in the Frobenius norm.

Oseledets and Tyrtyshnikov [OT10] generalize pseudoskeletons to tensors using the TT decomposition. Roughly, the idea for their algorithm, called *TT-cross*, is to replace the SVD in TT-SVD with pseudoskeleton decomposition. Importantly, for this to be efficient we should not always compute all of the factors in the pseudoskeleton. For example, if we are decomposing “from left to right”, i.e. first $k=1$, then $k=2$, etc, then we should only be computing the first two terms in (2.61) until $k=d$ where we compute the full pseudoskeleton. This allows for a very efficient implementation. Because we have not discussed submatrix selection here, we do not go into the details of TT-cross' computational complexity, but just mention that it is *linear* in d rather than exponential like naive TT-SVD. Moreover, Savostyanov [Sav14] shows that if the pseudoskeleton is chosen optimally in each step, then TT-cross is quasi-optimal in the sense that its error is a factor times the error of the best approximating tensor train.

If I was stranded on a desert island and could only bring one general-purpose tensor decomposition method, I think I would bring TT-cross.

Several software libraries implement TT decomposition algorithms. `TensorLy` implements TT-SVD. Python library `teneva` implements TT-ALS, TT-SVD, and TT-cross. Matlab library `TT-Toolbox` implements a lot of decompositions, including TT-ALS, TT-SVD, and TT-cross. Julia library `TensorToolbox.jl` implements TT-SVD.

2.5.6 Other decompositions

We have mentioned the CP, Tucker, and TT decompositions. There are many more tensor decompositions, but these are the focus of chapter 3. For examples of other tensor decompositions, one rich source of inspiration is Penrose diagrams. The term *tensor network* is often used. One decomposition that we did not discuss here, but is very useful, is *Hierarchical Tucker decomposition*. It is defined as a decomposition whose Penrose diagram is a binary tree. We also mention the survey by Grasedyck, Kressner, and Tobler [GKT13].

2.6 Appendix: proof of proposition 2.19

Proof. For each $x \in G$, we have a decomposition

$$T_x G = H_x \oplus V_x, \quad (2.63)$$

where $H_x = (\ker d\pi)^\perp$ and $V_x = \ker d\pi$ are the horizontal and vertical spaces respectively. We want to show that the metric

$$\langle d\pi_x r, d\pi_x s \rangle = \langle r, s \rangle, \quad r, s \in H_x \quad (2.64)$$

on G/H makes π into a Riemannian submersion. It is clear that the inner product is preserved as long as it is well-defined. So let $x' \in G$ such that $\pi(x) = \pi(x')$. Let $r', s' \in H_{x'}$ such that $d\pi_x r = d\pi_{x'} r'$ and $d\pi_x s = d\pi_{x'} s'$. Let R_x denote right translation by x . Then, since $d\pi_x$ is invertible on H_x ,

$$\langle s, r \rangle = \langle d\pi_x^{-1} d\pi_x r, d\pi_x^{-1} d\pi_x s \rangle \quad (2.65)$$

$$= \langle d\pi_x^{-1} d\pi_{x'} r', \dots \rangle \quad (2.66)$$

$$= \langle d(R_{x^{-1}x'})_x d\pi_x^{-1} d\pi_{x'} r', \dots \rangle \quad (2.67)$$

$$= \left\langle d(R_{x'^{-1}x})_{x'}^{-1} d\pi_x^{-1} d\pi_{x'} r', \dots \right\rangle \quad (2.68)$$

$$= \left\langle d(\pi \circ R_{xx'^{-1}})_{x'}^{-1} d\pi_{x'} r', \dots \right\rangle \quad (2.69)$$

$$\{xx'^{-1} \in H \implies \pi \circ R_{xx'^{-1}x} = \pi\} \quad (2.70)$$

$$= \langle d\pi_{x'}^{-1} d\pi_{x'} r', \dots \rangle \quad (2.71)$$

$$= \langle r', s' \rangle. \quad (2.72)$$

□

3. *Low-rank tensor manifolds*

This chapter is mostly based on arXiv preprint 2512.13594 “A homogeneous geometry of low-rank tensors” [Jac25], where I am the only author.

3.1 Introduction

If a rank- r matrix is multiplied from the left and right by invertible square matrices, then the product is again rank r . Moreover, any other rank- r matrix can be reached this way. In Lie algebraic terms, there is a *transitive action* G on the manifold of rank- r matrices.

A rank- r decomposition is unique up to an r -dimensional change of basis. Thus the group of such basis changes, together with the subgroup that fixes the rank factors, determines the *stabilizer* H , which is the subgroup of G that fixes a given rank- r matrix. In this way, the manifold of rank- r matrices is identified with the *homogeneous space* G/H .

A natural question is whether a similar construction also works for tensors. While there is more than one way to define the *rank* of a tensor, there are in general two obstacles to such a construction. First, the action G defined by multiplying in each mode by an invertible matrix might not be transitive. For example, there is no way to go from expressions like $a \otimes a \otimes a + b \otimes b \otimes b$ to $a \otimes a \otimes a + a \otimes b \otimes b$. Second, the uniqueness of the corresponding decomposition is more involved. Tensor decompositions are often studied from an algebraic geometric perspective, where decompositions are points on so-called *secant varieties*. Those varieties are often considered case-by-case, and broader results have to assume some symmetry, small size, or low rank. See for example [COV14; BC24].

In this paper, we consider three different ranks: *canonical polyadic* (CP) rank, *multilinear* rank, and *tensor train* (TT) rank. We give sufficient conditions

for when the set of fixed rank tensors is a homogeneous space under G . In those cases, fixed rank tensors form a smooth manifold and there is a canonical Riemannian metric defined by left-translating the Euclidean metric on G 's Lie algebra. We also show how the geodesics can be computed efficiently.

In low-rank approximation, the manifold perspective has been used to study several different fixed rank tensor spaces. We mention here CP rank [BV18], TT rank [HRS12a; UV20], hierarchical Tucker rank [UV13], and multilinear rank [KL10]. We also mention that fixed rank tensor manifolds have, to our knowledge, not previously been described as a homogeneous space, nor equipped with a Riemannian metric with known geodesics, other than in the rank 1 case [SvV22; Jac+24].

Identifying manifolds as homogeneous can have very profitable applications. For example, in *Riemannian optimization*, line searches and gradient descents are often performed along geodesics (or along approximate geodesics, called retractions). Using the canonical left-translated metric on a homogeneous manifold has two big advantages: geodesics on a quotient are easily described in terms of geodesics on its numerator and those geodesics are always complete, meaning they can be extended to arbitrary length. As mentioned, the space of matrices with fixed rank is a homogeneous manifold [VAV12; AO15; MV15], with natural applications to low-rank approximation.

Another application where homogeneity is relevant is *Lie group integrators*, a class of numerical integrators for ordinary differential equations (ODEs) on manifolds [Ise+00; CMO11; Bla+09]. They work by replacing the manifold ODE with an appropriate Lie group ODE, and identifying the underlying manifold as homogeneous informs the choice of Lie group and the design of the integrator [MZ97; Mun99; CO03; MW08; MQZ14; MV15].

The main novel contribution in this chapter is the identifications in theorems 3.3, 3.8, and 3.13 of sets of fixed-rank tensors as homogeneous manifolds. While some of these sets have previously been identified as smooth manifolds, their homogeneous structure have not previously been elucidated. We moreover classify in propositions 3.4, 3.9, and 3.14 when the homogeneous structure is reductive. Lastly, the expressions for the geodesics are also novel, and the results in propositions 3.5, 3.10, and 3.15 that these can be computed efficiently.

We now make precise what the action G is that we will use in this chapter. Let $d \geq 3$ and let n_1, \dots, n_d be integers ≥ 2 . If $\text{GL}(n)$ denotes the real group

of $n \times n$ invertible matrices, then

$$G = \mathrm{GL}(n_1) \times \cdots \times \mathrm{GL}(n_d) \quad (3.1)$$

has a natural action on the set of tensors $\mathbb{R}^{n_1} \otimes \cdots \otimes \mathbb{R}^{n_d}$ via

$$(g_1, \dots, g_d) \cdot (v_1 \otimes \cdots \otimes v_d) = (g_1 v_1) \otimes \cdots \otimes (g_d v_d) \quad (3.2)$$

and extended linearly to higher ranks. This is the Lie group action we will use to construct our manifolds.

We will write $G \curvearrowright M$ for “the action of G on M ”.

3.2 The CP manifold

Recall Kruskal’s theorem from section 2.5.3 or from Landsberg [Lan12, Theorem 12.5.3.2]. Landsberg formulates Kruskal’s theorem for complex vector spaces, but if a decomposition is real and complex unique then it is of course also real unique. The condition (2.37) is called *Kruskal’s criterion*. Whenever we use Kruskal’s theorem in the paper, we are using it via the following corollary.

Corollary 1. *Let $T \in \mathbb{R}^{n_1} \otimes \cdots \otimes \mathbb{R}^{n_d}$ be CP rank r with CP decomposition (2.35). Assume that, for each mode i , the factors v_i^1, \dots, v_i^r are linearly independent. Then the decomposition is unique up to permutation of the terms and rescaling of the factors.*

Note that v_i^1, \dots, v_i^r linearly independent in \mathbb{R}^{n_i} implies $r \leq n_i$ for all i .

3.2.1 Smooth manifold

Let $\mathrm{rank}^{-1}(r)$ denote the set of tensors with CP rank r . Its closure, $\overline{\mathrm{rank}^{-1}(r)}$, is an algebraic variety.

Lemma 3.1. *Let $r \leq n_i$ for all i . Then*

$$G \curvearrowright \mathrm{rank}^{-1}(r) \quad (3.3)$$

has an open dense orbit, Σ_r , consisting of elements with maximal multilinear rank.

Remark. Maximal multilinear rank in $\mathrm{rank}^{-1}(r)$ is (r, \dots, r) .

Proof. CP rank is preserved by the action of G since

1. if $\sum_{j=1}^r v_1^j \otimes \cdots \otimes v_d^j$ is a rank- r decomposition of T , then $\sum_{j=1}^r (g_1 v_1^j) \otimes \cdots \otimes (g_d v_d^j)$ is a rank- r decomposition of $g \cdot T$, so $\text{rank}(T) \geq \text{rank}(g \cdot T)$, and
2. if $\sum_{j=1}^r v_1^j \otimes \cdots \otimes v_d^j$ is a rank- r decomposition of $g \cdot T$, then $\sum_{j=1}^r (g_1^{-1} v_1^j) \otimes \cdots \otimes (g_d^{-1} v_d^j)$ is a rank- r decomposition of T , so $\text{rank}(g \cdot T) \geq \text{rank}(T)$.

Thus the action is well-defined.

Let e_i^1, \dots, e_i^d be the standard basis for \mathbb{R}^{n_i} and define

$$T = \sum_{j=1}^r e_1^j \otimes \cdots \otimes e_d^j. \quad (3.4)$$

We now want to show that $\Sigma_r = G \cdot T$ is Zariski open in $\overline{\text{rank}^{-1}(r)}$. It then follows that Σ_r is open and dense in $\text{rank}^{-1}(r)$.

Any other tensor

$$S = \sum_{j=1}^r f_1^j \otimes \cdots \otimes f_d^j. \quad (3.5)$$

with linearly independent f_i^1, \dots, f_i^r is reached by the group element (g_1, \dots, g_d) satisfying $g_i e_i^j = f_i^j$. Moreover, if f_i^1, \dots, f_i^r are linearly dependent for some i , S can't be reached by G . The complement of Σ_r is hence described by a Zariski closed condition, so Σ_r is Zariski open.

To show that elements in Σ_r have maximal multilinear rank, note that the matrix

$$T_{(i)} = \sum_{j=1}^r (e_1^j \otimes \cdots \otimes \widehat{e_i^j} \otimes \cdots \otimes e_d^j) \otimes e_i^j \quad (3.6)$$

has rank r since $e_1^j \otimes \cdots \otimes \widehat{e_i^j} \otimes \cdots \otimes e_d^j$ are r linearly independent vectors. Moreover, the action of G preserves multilinear rank.

To show that elements outside of Σ_r do not have maximal multilinear rank, note that if f_i^1, \dots, f_i^r are linearly dependent for some i , and

$$R = \sum_{j=1}^n f_1^j \otimes \cdots \otimes f_d^j, \quad (3.7)$$

then the column space of $R_{(i)}$ is spanned by $r - 1$ vectors. \square

We now consider the subgroup of G that fixes the T defined in (3.4). Applying Kruskal's theorem, we have the following result.

Lemma 3.2. *The stabilizer H of $G \curvearrowright \Sigma_r$ consists of elements of the form*

$$\begin{bmatrix} D_1 Q & M_1 \\ & A_1 \end{bmatrix} \times \cdots \times \begin{bmatrix} D_d Q & M_d \\ & A_d \end{bmatrix} \quad (3.8)$$

where D_i are invertible $r \times r$ diagonal matrices such that $D_1 \cdots D_d = I$, Q is an $r \times r$ permutation matrix, A_i are invertible $(n_i - r) \times (n_i - r)$ matrices, and M_i are $r \times (n_i - r)$ matrices.

Combining lemmas 3.1 and 3.2, we have the main result of this subsection.

Theorem 3.3. *The set of tensors with CP rank r and multilinear rank (r, \dots, r) is a smooth homogeneous manifold,*

$$\Sigma_r = G/H. \quad (3.9)$$

3.2.2 Problems with multilinear ranks other than (r, \dots, r)

One might imagine that a version of theorem 3.3 is true with (r, \dots, r) replaced by something smaller. However, such a statement is not true in general. The problem is that G is not transitive. Let a , b , and c be linearly independent vectors, and consider the two tensors

$$T = a \otimes a \otimes a + b \otimes b \otimes b + c \otimes c \otimes b \quad (3.10)$$

$$S = a \otimes a \otimes a + b \otimes b \otimes b + c \otimes c \otimes (a + b). \quad (3.11)$$

Both T and S have CP rank 3 (CP rank lower bounded by maximum of multilinear ranks) and multilinear rank $(3, 3, 2)$, yet there is no element in G that takes T to S . Here is a motivation: Kruskal's theorem applies to S , so if there was an action taking T to S then it would map a to a , b to b , and b to $a + b$, which is a contradiction. Note moreover that Kruskal's theorem does not apply to T .

On top of these complications, solving for the stabilizer is more difficult when Kruskal's theorem does not apply. Therefore, we settle for describing the situation when the multilinear rank is full.

3.2.3 Representatives

A point $p = gH \in \Sigma_r$ can be defined by specifying $g = (g_1, \dots, g_d)$. However, this requires specifying $n_1^2 + \cdots + n_d^2$ numbers, while the dimension of Σ_r is only

$(n_1 + \dots + n_d)r - (d-1)r$. To address this, we observe that a generic $g_i \in \text{GL}(n_i)$ can be reduced to block lower triangular form by the action of H . Define h_i via

$$g_i = \begin{bmatrix} g_{11} & g_{12} \\ g_{21} & g_{22} \end{bmatrix} = \begin{bmatrix} g_{11} & \\ g_{21} & 1 \end{bmatrix} \underbrace{\begin{bmatrix} 1 & g_{11}^{-1} g_{12} \\ g_{22} - g_{21} g_{11}^{-1} g_{12} \end{bmatrix}}_{h_i^{-1}}, \quad (3.12)$$

and note that $h = (h_1, \dots, h_d) \in H$. Here, we see that g_{11} needs to be invertible. If g_{11} is not invertible, there is a permutation matrix P such that $g'_i = P g_i$ has an invertible block g'_{11} . In this way we only need to specify g_{11} , g_{21} , and P for each i .

In practice, we want to choose P so that the determinant of g_{11} is maximized. This is known as the *submatrix selection* problem. It has been studied in depth because of its application to CUR-type matrix decompositions. See for example the review paper by Halko, Martinsson, and Tropp [HMT11]. We also mention the recent paper by Osinsky [Osi25] showing that quasioptimal choice can be computed in $\mathcal{O}(nr^2)$ basic operations, and the randomized extension proposed by Cortinovis and Kressner [CK25]. Hence there are efficient algorithms to choose g_{11} .

3.2.4 Riemannian manifold

G 's Lie algebra, \mathfrak{g} , consists of elements $Z = (Z_1, \dots, Z_d)$ where the Z_i are arbitrary $n_i \times n_i$ matrices. We consider the Euclidean inner product on \mathfrak{g} , defined as

$$\langle Z, Z' \rangle = \text{tr}(Z_1 Z_1'^T) + \dots + \text{tr}(Z_d Z_d'^T). \quad (3.13)$$

H 's Lie algebra, \mathfrak{h} , is a Lie subalgebra of \mathfrak{g} and by taking the derivative of the expression (3.8) in lemma 3.2 we see that \mathfrak{h} consists of elements $Y = (Y_1, \dots, Y_d)$ where the Y_i are on block form $\begin{bmatrix} Y_{11}^i & Y_{12}^i \\ & Y_{22}^i \end{bmatrix}$ with Y_{11}^i diagonal such that $Y_{11}^1 + \dots + Y_{11}^d = 0$, and Y_{12}^i and Y_{22}^i arbitrary. The orthogonal complement of \mathfrak{h} , which we denote \mathfrak{m} , thus consists of elements $X = (X_1, \dots, X_d)$ where the X_i are on block form $\begin{bmatrix} X_{11}^i & \\ X_{21}^i & 0 \end{bmatrix}$ such that X_{11}^i has the same diagonal for every i but is otherwise arbitrary, and X_{21}^i is arbitrary.

For any $g \in G$, the coset gH is a submanifold of G . Its tangent space at g is called the *vertical space*, and is denoted \mathcal{V}_g . The orthogonal complement to

the vertical space is called *horizontal space*, and is denoted \mathcal{H}_g . Thus \mathfrak{h} and \mathfrak{m} are the vertical and horizontal spaces at 1.

Now, consider the right-invariant metric¹ on G induced by the inner product on \mathfrak{g} . The *canonical metric* on Σ_r is then defined by demanding that the quotient map $\pi: G \rightarrow \Sigma_r$ is a Riemannian submersion. By construction, $d\pi_g|_{\mathcal{H}_g}: \mathcal{H}_g \rightarrow T_{gH}\Sigma_r$ is a linear isomorphism. So in other words, we are defining our metric by demanding that $d\pi_g|_{\mathcal{H}_g}$ is also an isometry.

Note also that the right-invariant metric on $GL(n)$ is left- $O(n)$ -invariant. In light of our discussion in section 3.2.3, this is important because permutation matrices are orthogonal and so we can work with any representative.

We now want to derive a more explicit description of the horizontal space. First, note that $\mathcal{V}_g = g\mathfrak{h}$, so that the vertical space consists of elements $Y = (Y_1, \dots, Y_d)$ where the $Y_i = g_i \begin{bmatrix} Y_{11} & Y_{12} \\ Y_{21} & Y_{22} \end{bmatrix}$. From now on, we suppress the i superscript, but note that $Y_{11} = Y_{11}^i$ and the other submatrices still depends on i . Second, if we choose a block lower triangular representative as in section 3.2.3, we have

$$\begin{aligned} & \langle X_i, Y_i \rangle_{g_i} \\ &= \text{tr} \begin{bmatrix} X_{11} & X_{12} \\ X_{21} & X_{22} \end{bmatrix} \begin{bmatrix} g_{11} & \\ g_{21} & 1 \end{bmatrix}^{-1} \left(\begin{bmatrix} g_{11} & \\ g_{21} & 1 \end{bmatrix} \begin{bmatrix} Y_{11} & Y_{12} \\ Y_{21} & Y_{22} \end{bmatrix} \begin{bmatrix} g_{11} & \\ g_{21} & 1 \end{bmatrix}^{-1} \right)^{\text{T}} \end{aligned} \quad (3.14)$$

$$= \text{tr} \begin{bmatrix} X_{11} & X_{12} \\ X_{21} & X_{22} \end{bmatrix} \begin{bmatrix} g_{11}^{-1} g_{11}^{-\text{T}} & -g_{11}^{-1} g_{11}^{-\text{T}} g_{21}^{\text{T}} \\ -g_{21} g_{11}^{-1} g_{11}^{-\text{T}} & 1 + g_{21} g_{11}^{-1} g_{11}^{-\text{T}} g_{21}^{\text{T}} \end{bmatrix} \begin{bmatrix} g_{11} Y_{11} & g_{11} Y_{12} \\ g_{21} Y_{11} & g_{21} Y_{12} + Y_{22} \end{bmatrix}^{\text{T}} \quad (3.15)$$

Since the Y_{12} is arbitrary, $Y'_{12} = g_{11} Y_{12}$ is also arbitrary. Similarly, $Y'_{22} = g_{21} Y_{12} + Y_{22}$ is arbitrary. Collecting the Y'_{12} coefficient yields the condition

$$-X_{11} g_{11}^{-1} g_{11}^{-\text{T}} g_{21}^{\text{T}} + X_{12} (1 + g_{21} g_{11}^{-1} g_{11}^{-\text{T}} g_{21}^{\text{T}}) = 0. \quad (3.16)$$

Note that $(1 + g_{21} g_{11}^{-1} g_{11}^{-\text{T}} g_{21}^{\text{T}})$ is invertible since it is the sum of a positive definite and positive semidefinite matrix, and so it is positive definite. We can hence solve for X_{12} :

$$X_{12} = X_{11} \Gamma_{12}, \quad (3.17)$$

¹We need the right-invariant metric rather than the left-invariant because we are dividing G by H on the right, so the metric on G needs to be at least right- H -invariant.

where $\Gamma_{12} = g_{11}^{-1} g_{11}^{-\top} g_{21}^{\top} (1 + g_{21} g_{11}^{-1} g_{11}^{-\top} g_{21}^{\top})^{-1}$. Similarly, collecting the Y'_{22} coefficients and solving for X_{22} yields

$$X_{22} = X_{21} \Gamma_{12}. \quad (3.18)$$

We also note that this implies that, for each i ,

$$X_i = \begin{bmatrix} X_{11} & X_{11} \Gamma_{12} \\ X_{21} & X_{21} \Gamma_{12} \end{bmatrix} = \begin{bmatrix} X_{11} \\ X_{21} \end{bmatrix} \begin{bmatrix} 1 & \Gamma_{12} \end{bmatrix} \quad (3.19)$$

is at most rank r .

The power series trick Before collecting the Y_{11} coefficients, we discuss an alternative expression for Γ_{12} that will allow us to compute it more efficiently. As it is currently written, building Γ_{12} requires inverting the $n_i \times n_i$ matrix $1 + g_{21} g_{11}^{-1} g_{11}^{-\top} g_{21}^{\top}$. Note however that this inverse is an analytic function of the rank r matrix $g_{21} g_{11}^{-1} g_{11}^{-\top} g_{21}^{\top}$. Its power series is

$$\frac{1}{1+x} = 1 - x + x^2 - x^3 + \dots \quad (3.20)$$

If $M = AB$ is a rank r decomposition of the $n \times n$ matrix M , then

$$\frac{1}{1+M} = 1 - AB + (AB)^2 - (AB)^3 + \dots \quad (3.21)$$

$$= 1 + A(-1 + BA - (BA)^2 + \dots)B \quad (3.22)$$

$$= 1 - A \frac{1}{1+BA} B. \quad (3.23)$$

But BA is an $r \times r$ matrix, allowing $-1/(1+BA)$ to be evaluated efficiently. This trick is also useful later for evaluating the matrix exponential. For Γ_{12} , we have the following expression,

$$\Gamma_{12} = g_{11}^{-1} \left(1 - \frac{[g_{11}^{-\top} g_{21}^{\top} g_{21} g_{11}^{-1}]}{1 + [g_{11}^{-\top} g_{21}^{\top} g_{21} g_{11}^{-1}]} \right) g_{11}^{-\top} g_{21}^{\top}. \quad (3.24)$$

Especially, the expression in brackets is an $r \times r$ matrix.

We are now ready to collect the Y_{11} coefficients. We get the condition that the k th column of

$$\begin{bmatrix} X_{11} \\ X_{21} \end{bmatrix} g_{11}^{-1} g_{11}^{-\top} - \begin{bmatrix} X_{12} \\ X_{22} \end{bmatrix} g_{21} g_{11}^{-1} g_{11}^{-\top} = \begin{bmatrix} X_{11} \\ X_{21} \end{bmatrix} (g_{11}^{-1} g_{11}^{-\top} - \Gamma_{12} g_{21} g_{11}^{-1} g_{11}^{-\top}) \quad (3.25)$$

has the same dot product with the k th column of $\begin{bmatrix} g_{11} \\ g_{12} \end{bmatrix}$ for every i .

This completes the description of the horizontal space \mathcal{H}_g . The most important takeaway is that horizontal vectors have a decomposition (3.19).

3.2.5 Riemannian homogeneous manifold

The construction described in section 3.2.4 uses a right-invariant metric on G , but the resulting metric on Σ_r is not necessarily right-invariant. In fact, not all smooth homogeneous manifolds allow for an invariant metric. The underlying issue is that since $h \cdot p = p$ for all $h \in H$, the inner product on $T_p \Sigma_r$ has to be invariant under $H \curvearrowright \Sigma_r$, but such an inner product might not exist.

The technical condition that we want to satisfy is that there exists a subspace $\mathfrak{p} \subset \mathfrak{g}$, with $\mathfrak{p} \oplus \mathfrak{h} = \mathfrak{g}$, such that $h\mathfrak{p}h^{-1} \subset \mathfrak{p}$ for all $h \in H$. G/H is then called *reductive* and \mathfrak{p} is called *invariant subspace*. See O’Neill [ONe83, Chapter 11] for more details.

Proposition 3.4. Σ_r is reductive if and only if $r = n_1 = \dots = n_d$. Moreover, when Σ_r is reductive, the \mathfrak{m} that we defined in section 3.2.4 is an invariant subspace.

Proposition 3.4 is a higher order version of [VAV12, Proposition 3.4] and [MV15, Proposition 5.7], which say that fixed rank matrices are only reductive when they are square and full rank. The parts of their proofs that deal with the “only if” direction can be recycled directly to prove proposition 3.4. However, the statement that our manifold is reductive depends on the precise form of H , and so here we need a specialized argument.

Proof. Assume $r = n_1 = \dots = n_d$ and let $h = D_1 Q \times \dots \times D_d Q \in H$. Let $X \in \mathfrak{m}$ and consider the expression

$$hXh^{-1} = (D_1 Q X_1 Q^{-1} D_1^{-1}, \dots, D_d Q X_d Q^{-1} D_d^{-1}). \quad (3.26)$$

If σ denotes the permutation that corresponds to Q , then on the diagonals we have that

$$\begin{aligned} (D_i Q X_i Q^{-1} D_i^{-1})_{kk} &= (D_i)_{kk} (Q X_i Q^{-1})_{kk} (D_i^{-1})_{kk} \\ &= (D_i)_{kk} (X_i)_{\sigma(k)\sigma(k)} (D_i^{-1})_{kk} \\ &= (X_i)_{\sigma(k)\sigma(k)}. \end{aligned} \quad (3.27)$$

Hence hXh^{-1} is contained in \mathfrak{m} , which is what we wanted to show.

Conversely, assume there exists an invariant subspace \mathfrak{p} when $n_i > r$ for some i . Note that elements in \mathfrak{p} are determined by their projection to \mathfrak{m} . Concretely, elements in \mathfrak{p} are determined by their r first columns. So consider an element $X = (X_1, \dots, X_d) \in \mathfrak{p}$ with

$$X_i = \begin{bmatrix} X_{11} & X_{12} \\ X_{21} & X_{22} \end{bmatrix}. \quad (3.28)$$

X_{12} and X_{22} are functions of X_{11} and X_{21} . For our contradiction, we will show that X_{11} is a multiple of the identity. Then the dimension of \mathfrak{p} is too low to be complementary to \mathfrak{h} .

First, assume $X_{21} = 0$ and let $h = \begin{bmatrix} 1 & M \\ & 1 \end{bmatrix}$ with M arbitrary. Then $h^{-1} = \begin{bmatrix} 1 & -M \\ & 1 \end{bmatrix}$ and

$$hX_ih^{-1} = \begin{bmatrix} X_{11} & -X_{11}M + X_{12} + MX_{22} \\ & X_{22} \end{bmatrix} \quad (3.29)$$

$$h^{-1}X_ih = \begin{bmatrix} X_{11} & X_{11}M + X_{12} - MX_{22} \\ & X_{22} \end{bmatrix}. \quad (3.30)$$

The first r columns of hX_ih^{-1} and $h^{-1}X_ih$ are the same, so by our previous comment they must be equal. Thus

$$X_{11}M - MX_{22} = 0 \quad (3.31)$$

for all M . Writing this as

$$(X_{11} \otimes 1 - 1 \otimes X_{22}) \text{vec } M = 0, \quad (3.32)$$

it is clear that the only solutions are when X_{11} and X_{22} are multiples of identity matrices and have the same norm.

Second, if $X_{21} \neq 0$, we can use the same argument on

$$\begin{bmatrix} X_{11} & X_{12} \\ X_{21} & X_{22} \end{bmatrix} - \begin{bmatrix} X'_{12} & X'_{11} \\ X'_{22} & X'_{21} \end{bmatrix} \in \mathfrak{p} \quad (3.33)$$

to see that X_{11} again is a multiple of the identity.

Whenever $r \geq 2$, the above is enough for a contradiction. However, when $r = 1$, X_{11} is just a 1×1 matrix, and so showing that it is a multiple of the

identity yields no contradiction. We need to change the argument slightly. Consider now instead the case $X_{11} = 0$.

$$hX_i h^{-1} = \begin{bmatrix} MX_{21} & * \\ X_{21} & X_{22} - X_{21}M \end{bmatrix}. \tag{3.34}$$

By the previous paragraph, we are free to add and subtract any multiple of the identity matrix and still stay in \mathfrak{p} . If we subtract the real number MX_{21} from the diagonal, we are left with

$$\begin{bmatrix} & * \\ X_{21} & X_{22} - X_{21}M - MX_{21} \cdot 1 \end{bmatrix}. \tag{3.35}$$

Similarly to before, since the first column is the same as X_i , it must be equal to X_i , but for this to be true for all M implies $X_{21} = 0$. X_{21} is ours to choose, so any restriction on it is a contradiction. \square

3.2.6 Geodesics

Since the subgroup associated with Q is discrete, and since being geodesic is a local property, we may for the purposes of this subsection ignore it and set $Q = 1$.

Let $g_i \in \text{GL}(n_i)$ and $X_i \in T_{g_i} \text{GL}(n_i)$. Andruchow, Larotonda, Recht, and Varela [And+14] show that the geodesics on the general linear group are²

$$\exp_{g_i}(X_i) = \text{mexp}(X_i g_i^{-1} - (X_i g_i^{-1})^\top) \text{mexp}((X_i g_i^{-1})^\top) g_i. \tag{3.36}$$

Geodesics on $\Sigma_r = G/H$ are images of horizontal geodesics in G under the quotient map $\pi: G \rightarrow G/H$. In this subsection, assuming $r \ll n_i$, our aim is to efficiently compute (an element in the same equivalence class as) (3.36) when X_i is part of a horizontal vector. We are going to show that this can be done without any $\mathcal{O}(n_i^2)$ operations, as would be required by lifting to G .

First, recall that horizontal vectors are on the form (3.19). We thus have a rank- r decomposition

$$X_i g_i^{-1} = \begin{bmatrix} X_{11} \\ X_{21} \end{bmatrix} [g_{11}^{-1} - \Gamma_{12} g_{21} g_{11}^{-1} \quad \Gamma_{12}]. \tag{3.37}$$

To compute the matrix exponential of such a vector, we can use the same power series trick as before. If we name the factors in (3.37) $A \in \mathbb{R}^{n_i \times r}$ and $B \in \mathbb{R}^{r \times n_i}$

²Note that they use a left-invariant metric while we use a right-invariant.

respectively, then

$$\begin{aligned} \text{mexp}(X_i g_i^{-1}) &= 1 + AB + \frac{1}{2}(AB)^2 + \dots \\ &= 1 + A\psi_1(BA)B, \end{aligned} \quad (3.38)$$

where $\psi_1(x) = 1 + \frac{1}{2}x + \frac{1}{3!}x^2 + \dots$ is the Taylor series for $(e^x - 1)/x$. The notation ψ_1 comes from the theory of exponential integrators, where the functions

$$\psi_k(x) = \sum_{j=i}^{\infty} \frac{1}{(j+k)!} x^j \quad (3.39)$$

are known as *the ψ functions*. Importantly, the argument to ψ_1 is an $r \times r$ matrix, so it can be evaluated cheaply. We will discuss exactly how shortly.

Second, we have a rank- $2r$ decomposition

$$X_i g_i^{-1} - (X_i g_i^{-1})^\top = \begin{bmatrix} X_{11} & -(g_{11}^{-1} - \Gamma_{12} g_{21} g_{11}^{-1})^\top \\ X_{21} & -\Gamma_{12}^\top \end{bmatrix} \begin{bmatrix} g_{11}^{-1} - \Gamma_{12} g_{21} g_{11}^{-1} & \Gamma_{12} \\ X_{11}^\top & X_{21}^\top \end{bmatrix}. \quad (3.40)$$

So, similarly to before, denoting the factors $A' \in \mathbb{R}^{n_i \times 2r}$ and $B' \in \mathbb{R}^{2r \times n_i}$ respectively,

$$\text{mexp}(X_i g_i^{-1} - (X_i g_i^{-1})^\top) = 1 + A'\psi_1(B'A')B'. \quad (3.41)$$

Here, the argument to ψ_1 is an $2r \times 2r$ matrix, so it can be evaluated cheaply.

Putting all this into (3.36) and doing the multiplications, we find an expression for the first r columns,

$$\begin{aligned} &\exp_{g_i}(X_i) \begin{bmatrix} 1 \\ 0 \end{bmatrix} \\ &= \begin{bmatrix} g_{11} \\ g_{21} \end{bmatrix} + A\psi_1(BA)B \begin{bmatrix} g_{11} \\ g_{21} \end{bmatrix} + A'\psi_1(B'A')B' \begin{bmatrix} g_{11} \\ g_{21} \end{bmatrix} \\ &\quad + A\psi_1(BA)BA'\psi_1(B'A')B' \begin{bmatrix} g_{11} \\ g_{21} \end{bmatrix} \end{aligned} \quad (3.42)$$

Advantageously, if the multiplications are done in the right order this does not require forming any $n_i \times n_i$ matrices.

We now return to how to compute ψ_1 . We will do it similarly to how the matrix exponential is usually computed, using *Padé approximation* and

scaling and squaring. See Moler and van Loan [MV03, methods 2 and 3] and Higham [Hig08, sections 10.3 and 10.7.4]. Let M be an $r \times r$ matrix and assume first that $\|M\| \leq 1/2$. Then $\psi_1(M)$ is approximated to double precision from its degree (6, 6) Padé approximant [Hig08, theorem 10.31]

$$\begin{aligned} r_{66}(M) &= \frac{1 + M/26 + 5M^2/156 + M^3/858 + M^4/5720 + M^5/205920 + M^6/8648640}{1 - 6M/13 + 5M^2/52 - 5M^3/429 + M^4/1144 - M^5/25740 + M^6/1235520}. \end{aligned} \quad (3.43)$$

We prove this in lemma 3.16.

On the other hand, if $\|M\| > 1/2$, then we do not use the Padé approximant directly. Let z be an integer such that $\|M\|2^{-z} \leq 1/2$. Keeping in mind that these expressions are only shorthands for their Taylor series, we have

$$\begin{aligned} \psi_1(M) &= \frac{\text{mexp}(M) - 1}{M} \\ &= \frac{\text{mexp}(2^{-z}M)^{2^z} - 1}{M} \\ &= \frac{\text{mexp}(2^{-z}M) - 1}{M} (\text{mexp}(2^{-z}M) + 1) \cdots (\text{mexp}(2^{-1}M) + 1) \\ &= 2^{-z} \psi_1(2^{-z}M) (\text{mexp}(2^{-z}M) + 1) \cdots (\text{mexp}(2^{-1}M) + 1). \end{aligned} \quad (3.44)$$

In the second to last step, we used $x^2 - 1 = (x - 1)(x + 1)$ recursively. This scaling and squaring step is similar to the algorithm proposed by Hochbruck, Lubich, and Selhofer [HLS98].

Counting the operations We now count the number of basic operations required to evaluate (3.42). We use the same conventions as [Hig08, Table C.1], where a $(a \times b) \times (b \times c)$ matrix multiplication is $2abc$ basic operations, and a $(a \times b) \times (b \times c)$ matrix division is $8abc/3$ basic operations. Terms that are $\mathcal{O}(n_i r)$ or $\mathcal{O}(r^2)$, such as adding $r \times r$ matrices, are ignored.

Proposition 3.5. *Given a tangent vector $X \in T_p \Sigma_r$, $\exp_p(X)$ can be estimated³ using*

$$\sum_{i=1}^d \left[\frac{110}{3} n_i r^2 + (146 + 36z_i) r^3 + \mathcal{O}(n_i r + r^2) \right] \text{ basic operations} \quad (3.45)$$

³Meaning that everything is computed exactly except for ψ_1 , which is Padé approximated to within double precision.

where⁴ $z_i = \lceil \log_2 \|X_i g_i^{-1}\| \rceil + 2$.

See section 3.8 for the proof.

This can be viewed as a tensorial version of Vandereycken's et al. [VAV12] corollary 4.2. We also mention that the effects of rounding errors in the Padé approximant and squaring step are not completely understood [MV03], and so we do not do a full error analysis.

3.3 The Tucker manifold

Recall from proposition 2.26 that the Tucker decomposition is unique up to a change of basis in the contracted edges.

3.3.1 Smooth manifold

Let $\text{mrank}^{-1}(t_1, \dots, t_d)$ denote the set of tensors with multilinear rank (t_1, \dots, t_d) . Like $\text{rank}^{-1}(r)$, its closure is an algebraic variety.

Lemma 3.6. *Let $t_1 = t_2 \cdots t_d$. Then*

$$G \curvearrowright \Lambda_{t_1 \dots t_d} := \text{mrank}^{-1}(t_1, \dots, t_d) \quad (3.46)$$

is a transitive action.

Proof. Let E_i be the identity matrix $I_{n_i \times t_i}$ seen as an element of $\mathbb{R}^{n_i} \otimes \mathbb{R}^{t_i}$ and let I be the identity matrix $I_{t_1 \times t_1}$ seen as an element of $\mathbb{R}^{t_1} \otimes \cdots \otimes \mathbb{R}^{t_d}$. Then define

$$T_{k_1 \dots k_d} = \sum_{\alpha_1, \dots, \alpha_d} I_{\alpha_1 \dots \alpha_d} (E_1)_{k_1}^{\alpha_1} \cdots (E_d)_{k_d}^{\alpha_d}. \quad (3.47)$$

Now, fix C and G_1, \dots, G_i . By a previous comment, a Tucker decomposition with inner dimensions t_1, \dots, t_d has G_i with linearly independent columns. So let $\overline{G}_i: \mathbb{R}^{n_i \times (n_i - t_i)}$ be a basis completion to G_i . Moreover, the unfolding $C: \mathbb{R}^{t_1} \rightarrow \mathbb{R}^{t_2} \otimes \cdots \otimes \mathbb{R}^{t_d}$ is an invertible $t_1 \times t_1$ matrix. The tensor

$$S_{k_1 \dots k_d} = \sum_{\alpha_1, \dots, \alpha_d} C_{\alpha_1 \dots \alpha_d} (G_1)_{k_1}^{\alpha_1} \cdots (G_d)_{k_d}^{\alpha_d}. \quad (3.48)$$

is thus reached by $g = ([G_1 \ \overline{G}_1] \begin{bmatrix} C \\ 1 \end{bmatrix}, [G_2 \ \overline{G}_2], \dots, [G_d \ \overline{G}_d])$. \square

⁴This has to be the same norm as in section 3.7.

From proposition 2.26 we can directly compute the stabilizer.

Lemma 3.7. *The stabilizer H of $G \curvearrowright \Lambda_{t_1 \dots t_d}$ consists of elements of the form*

$$\begin{bmatrix} A_2^{-\top} \otimes \cdots \otimes A_d^{-\top} & M_1 \\ & B_1 \end{bmatrix} \times \begin{bmatrix} A_2 & M_2 \\ & B_2 \end{bmatrix} \times \cdots \times \begin{bmatrix} A_d & M_d \\ & B_d \end{bmatrix} \quad (3.49)$$

where the A_i are invertible $t_i \times t_i$ matrices, the B_i are invertible $(n_i - t_i) \times (n_i - t_i)$ matrices, and M_i are $t_i \times (n_i - t_i)$ matrices.

Theorem 3.8. *The set of tensors with multilinear rank (t_1, \dots, t_d) , where $t_1 = t_2 \cdots t_d$, is a smooth homogeneous manifold,*

$$\Lambda_{t_1 \dots t_d} = G/H. \quad (3.50)$$

3.3.2 Problems with multilinear ranks other than $t_1 = t_2 \cdots t_d$

There are other cases than $t_1 = t_2 \cdots t_d$ where $\text{mrnk}^{-1}(t_1, \dots, t_d)$ has an open and dense orbit. However, describing those orbits is more work and involves the so-called *castling transform* of the factors, which generalizes the statement that $\mathbb{R}^p \otimes \mathbb{R}^q \otimes \mathbb{R}^r$ has an open orbit iff $\mathbb{R}^p \otimes \mathbb{R}^q \otimes \mathbb{R}^{p+q-r}$ has an open orbit. See Venturelli [Ven18] or Landsberg [Lan12, Section 10.2.2] for details. We also repeat the observation from the introduction: that there is typically more degrees of freedom in $\mathbb{R}^{t_1} \otimes \cdots \otimes \mathbb{R}^{t_d}$, namely $t_1 \cdots t_d$ degrees, than in $\text{GL}(t_1) \times \cdots \times \text{GL}(t_d)$, namely $t_1^2 + \cdots + t_d^2$ degrees. This imposes the restriction

$$t_1 \cdots t_d \leq t_1^2 + \cdots + t_d^2. \quad (3.51)$$

Furthermore, since $t_1 > t_2 \cdots t_d$ is not a possible multilinear rank, we also have the restriction

$$t_1 \leq t_2 \cdots t_d. \quad (3.52)$$

Solving for t_1 , we find that it must satisfy

$$\frac{t_2 \cdots t_d + \sqrt{(t_2 \cdots t_d)^2 - 4(t_2^2 + \cdots + t_d^2)}}{2} \leq t_1 \leq t_2 \cdots t_d. \quad (3.53)$$

This domain is typically very small. For $d = 3$, $t_2 = t_3 = 10$ for example, we have $98 \leq t_1 \leq 100$. We therefore settle for describing the case $t_1 = t_2 \cdots t_d$.

3.3.3 Representatives

Similarly to section 3.2.3, if $g = (g_1, \dots, g_d) \in G$, then we only need to store the first t_i columns of g_i .

3.3.4 Riemannian manifold

Similarly to section 3.2.4, we can take the derivative of an arbitrary curve in (3.49) in lemma 3.7 to get H 's Lie algebra, \mathfrak{h} . It consists of elements $Y = (Y_1, \dots, Y_d)$ where the Y_i are on block form

$$\begin{aligned} Y_1 &= \begin{bmatrix} -K_2^\top \otimes 1 \otimes \cdots \otimes 1 - \cdots - 1 \otimes \cdots \otimes 1 \otimes K_d^\top & * \\ & * \end{bmatrix}, \\ Y_i &= \begin{bmatrix} K_i & * \\ & * \end{bmatrix}, \quad 2 \leq i \leq d, \end{aligned} \quad (3.54)$$

for arbitrary $t_i \times t_i$ matrices K_i .

The orthogonal complement, \mathfrak{m} , consists of elements $X = (X_1, \dots, X_d)$ where the X_i are on block form

$$\begin{aligned} X_1 &= \begin{bmatrix} L_1 & \\ * & 0 \end{bmatrix}, \\ X_i &= \begin{bmatrix} L_i & \\ * & 0 \end{bmatrix}, \quad 2 \leq i \leq d, \end{aligned} \quad (3.55)$$

such that

$$L_i = \text{tr}_i L_1, \quad 2 \leq i \leq d, \quad (3.56)$$

where $\text{tr}_i(A_1 \otimes \cdots \otimes A_d) = (\text{tr } A_1) \cdots \widehat{\text{tr } A_i} \cdots (\text{tr } A_d) A_i^\top$. Note that L_2, \dots, L_d are completely determined by L_1 .

Like in section 3.2.4, we consider the right-invariant metric on G induced by the Euclidean inner product on \mathfrak{g} , and define the metric on $\Lambda_{t_1 \dots t_d}$ by demanding that $\pi: G \rightarrow \Lambda_{t_1 \dots t_d}$ is a Riemannian submersion. We do not give a full description of the horizontal space at a general point, but just note that the same argument that was used in section 3.2.4 can be used to derive a rank- t_i decomposition similar to (3.19).

3.3.5 Riemannian homogeneous manifold

Proposition 3.9. $\Lambda_{t_1 \dots t_d}$ is reductive if and only if $n_1 = t_1, \dots, n_d = t_d$. Moreover, when $\Lambda_{t_1 \dots t_d}$ is reductive, \mathfrak{m} is an invariant subspace.

Proof. Assume $n_1 = t_1, \dots, n_d = t_d$ and let $h = ((A_2 \otimes \cdots \otimes A_d)^\top, A_2, \dots, A_d) \in H$. We have to show that (3.56) is preserved by $L_i \mapsto L'_i = A_i L_i A_i^{-1}$. This can

be seen from

$$\begin{aligned}
L'_i &= A_i L_i A_i^{-1} \\
&= A_i (\text{tr}_i L_1) A_i^{-1} \\
&= \text{tr}_i [(A_1 \otimes \cdots \otimes A_d)^{-\top} L_1 (A_1 \otimes \cdots \otimes A_d)^\top] \\
&= \text{tr}_i L'_1.
\end{aligned} \tag{3.57}$$

If $n_i \neq t_i$ for some i , then it is possible to show that $\Lambda_{t_1 \dots t_d}$ is not reductive with the same argument as in proposition 3.4. \square

3.3.6 Geodesics

By an argument completely analogous to proposition 3.5, we have the following result.

Proposition 3.10. *Given a tangent vector $X \in T_p \Lambda_{t_1 \dots t_d}$, $\exp_p(X)$ can be estimated using*

$$\sum_{i=1}^d \left[\frac{110}{3} n_i t_i^2 + (146 + 36z_i) t_i^3 + \mathcal{O}(n_i t_i + t_i^2) \right] \text{ basic operations} \tag{3.58}$$

where $z_i = \lceil \log_2 \|X_i g_i^{-1}\| \rceil + 2$.

3.4 The tensor train manifold

Recall Uschmajew and Vandereycken's uniqueness result from proposition 2.29 about TT decompositions.

Many of the arguments in this section are the same as in section 3.2, so we do not repeat them here but just refer back.

3.4.1 Smooth manifold

Let $\text{ttrank}^{-1}(s_1, \dots, s_{d-1})$ denote the set of tensors with TT rank (s_1, \dots, s_{d-1}) . Like $\text{rank}^{-1}(r)$, its closure is an algebraic variety.

Lemma 3.11. *Let $s_1 \leq n_1$, $s_{i-1}s_i \leq n_i$ for all $2 \leq i \leq d-1$, and $s_{d-1} \leq n_d$. Then*

$$G \curvearrowright \text{ttrank}^{-1}(s_1, \dots, s_{d-1}) \quad (3.59)$$

has an open dense orbit, $\Pi_{s_1 \dots s_{d-1}}$, consisting of elements with maximal multilinear rank.

Remark. Maximal multilinear rank in $\text{ttrank}^{-1}(s_1, \dots, s_{d-1})$ is $(s_1, s_1 s_2, \dots, s_{d-1})$.

Proof. From (2.50), it is clear that TT rank is preserved under G . The action is thus well-defined.

Let $E_i: \mathbb{R}^{s_{i-1}} \otimes \mathbb{R}^{n_i} \otimes \mathbb{R}^{s_i}$ be the tensor whose unfolding is the identity matrix, $(E_i: \mathbb{R}^{n_i} \rightarrow \mathbb{R}^{s_{i-1}} \otimes \mathbb{R}^{s_i}) = I_{n_i \times s_{i-1} s_i}$, and define

$$T_{k_1 k_2 \dots k_d} = \sum_{\alpha_1, \dots, \alpha_{d-1}} (E_1)_{k_1 \alpha_1} (E_2)^{\alpha_1 k_2 \alpha_2} \dots (E_d)^{\alpha_{d-1} k_d}. \quad (3.60)$$

Similarly to the proof of lemma 3.1, we now want to show that $\Pi_{s_1 \dots s_{d-1}} = G \cdot T$ is Zariski open in $\text{ttrank}^{-1}(s_1, \dots, s_{d-1})$.

Any other tensor

$$S_{k_1 \dots k_d} = \sum_{\alpha_1, \dots, \alpha_{d-1}} (F_1)_{k_1 \alpha_1} (F_2)^{\alpha_1 k_2 \alpha_2} \dots (F_d)^{\alpha_{d-1} k_d} \quad (3.61)$$

where the matrices $F_i: \mathbb{R}^{n_i \times s_{i-1} s_i}$ have full rank is reached by the group element (g_1, \dots, g_d) where $g_i = [F_i \ \bar{F}_i]$ such that $\bar{F}_i: \mathbb{R}^{n_i \times (n_i - s_{i-1} s_i)}$ is a basis completion. Moreover, the action of G preserves the rank of F_i . The condition that the F_i have full rank is a Zariski open condition.

To show that elements outside of $\Pi_{s_1 \dots s_{d-1}}$ do not have maximal multilinear rank, note that the multilinear rank of S is just $(\text{rank } F_1, \dots, \text{rank } F_d)$. \square

From Proposition 2.29 we can directly compute the stabilizer.

Lemma 3.12. *The stabilizer H of $G \curvearrowright \Pi_{s_1 \dots s_{d-1}}$ consists of elements of the form*

$$\begin{bmatrix} A_1 & M_1 \\ & B_1 \end{bmatrix} \times \begin{bmatrix} A_1^{-\top} \otimes A_2 & M_2 \\ & B_2 \end{bmatrix} \times \dots \times \begin{bmatrix} A_{d-1}^{-\top} & M_d \\ & B_d \end{bmatrix} \quad (3.62)$$

where the A_i are invertible $s_i \times s_i$ matrices, the B_i are invertible $(n_i - s_{i-1} s_i) \times (n_i - s_{i-1} s_i)$ matrices, and the M_i are $s_{i-1} s_i \times (n_i - s_{i-1} s_i)$ matrices.

Combining lemmas 3.11 and 3.12, we have the main result of this subsection.

Theorem 3.13. *The set of tensors with TT rank (s_1, \dots, s_{d-1}) and multilinear rank $(s_1, s_1 s_2, \dots, s_{d-1})$ is a smooth homogeneous manifold,*

$$\Pi_{s_1 \dots s_{d-1}} = G/H. \quad (3.63)$$

3.4.2 Problems with multilinear ranks other than $(s_1, s_1 s_2, \dots, s_d)$

Like theorem 3.3, one might imagine that a version of theorem 3.13 is true with $(s_1, s_1 s_2, \dots, s_{d-1})$ replaced by something smaller. However such a statement is not true in general. The problem is once again that G is not transitive. Consider a TT rank $(10, 10)$ tensor with multilinear rank $(10, 97, 10)$, i.e. slightly less than the maximal rank $(10, 100, 10)$. The dimension of G 's action is $10^2 + 97^2 + 10^2 = 9609$, while the dimension of the set of such TT tensors is at least $10 \cdot 10 \cdot 97 = 9700$. Thus, G can not be transitive. Recall from section 2.3.1 the definition of *transitive*.

Therefore, we settle for describing the situation when the multilinear rank is full. Compared to CP tensors, the computation of the stabilizer should be straightforward as everything should reduce to matrix-type uniqueness questions in each mode.

3.4.3 Representatives

Similarly to section 3.2.3, if $g = (g_1, \dots, g_d) \in G$, then we only need to store the first $s_{i-1} s_i$ columns of g_i .

3.4.4 Riemannian manifold

Similarly to section 3.2.4, we can take the derivative of the expression (3.62) in lemma 3.12 to get H 's Lie algebra, \mathfrak{h} . It consists of elements $Y = (Y_1, \dots, Y_d)$ where the Y_i are on block form

$$\begin{aligned} Y_1 &= \begin{bmatrix} K_1 & * \\ & * \end{bmatrix}, \\ Y_i &= \begin{bmatrix} -K_{i-1}^T \otimes 1 + 1 \otimes K_i & * \\ & * \end{bmatrix}, \quad 2 \leq i \leq d-1, \\ Y_d &= \begin{bmatrix} -K_{d-1}^T & * \\ & * \end{bmatrix}, \end{aligned} \quad (3.64)$$

for arbitrary $s_i \times s_i$ matrices K_i . The orthogonal complement, \mathfrak{m} , consists of elements $X = (X_1, \dots, X_d)$ where the X_i are on block form

$$\begin{aligned} X_1 &= \begin{bmatrix} L_1 & \\ * & 0 \end{bmatrix}, \\ X_i &= \begin{bmatrix} L_i & \\ * & 0 \end{bmatrix}, \quad 2 \leq i \leq d-1, \\ X_d &= \begin{bmatrix} L_d \\ * & 0 \end{bmatrix}, \end{aligned} \tag{3.65}$$

such that

$$\begin{aligned} L_1 &= \operatorname{tr}_2 L_2 \\ \operatorname{tr}_1 L_{i-1} &= \operatorname{tr}_2 L_i, \quad 2 \leq i \leq d-1 \\ \operatorname{tr}_1 L_{d-1} &= L_d, \end{aligned} \tag{3.66}$$

where $\operatorname{tr}_1(A \otimes B) = (\operatorname{tr} A)B$ and $\operatorname{tr}_2(A \otimes B) = (\operatorname{tr} B)A^\top$. The easiest way to see this restriction on L_i is to write the Euclidean inner product as

$$\langle Y, X \rangle = \langle K_1, L_1 - \operatorname{tr}_2 L_2 \rangle + \langle K_2, \operatorname{tr}_1 L_2 - \operatorname{tr}_2 L_3 \rangle + \dots + \langle K_{d-1}, \operatorname{tr}_1 L_{d-1} - L_d \rangle \tag{3.67}$$

and note that this should hold for all K_i separately.

Like in section 3.2.4, we consider the right-invariant metric on G induced by the Euclidean inner product on \mathfrak{g} , and define the metric on $\Pi_{s_1 \dots s_{d-1}}$ by demanding that $\pi: G \rightarrow \Pi_{s_1 \dots s_{d-1}}$ is a Riemannian submersion. We do not give a full description of the horizontal space at a general point g , but just note that the same argument that was used in section 3.2.4 can be used to derive a decomposition similar to (3.19) with inner dimension $s_{i-1}s_i$,

$$X_i = \begin{bmatrix} X_{11} \\ X_{21} \end{bmatrix} \begin{bmatrix} 1 & \Gamma_{12} \end{bmatrix}. \tag{3.68}$$

3.4.5 Riemannian homogeneous manifold

Proposition 3.14. $\Pi_{s_1 \dots s_{d-1}}$ is reductive if and only if $n_1 = s_1, n_2 = s_1 s_2, \dots, n_d = s_{d-1}$. Moreover, when $\Pi_{s_1 \dots s_{d-1}}$ is reductive, \mathfrak{m} is an invariant subspace.

Proof. Assume $n_1 = s_1, n_2 = s_1 s_2, \dots, n_d = s_{d-1}$ and let $h = (A_1, A_1^{-\top} \otimes A_2, \dots, A_{d-1}^{-\top}) \in H$. We have that $X \mapsto hXh^{-1}$ maps $L_i \mapsto L'_i = (A_{i-1}^{-\top} \otimes$

$A_i)L_i(A_{i-1}^\top \otimes A_i^{-1})$. To show that \mathfrak{m} is invariant, we need to argue that the relation (3.66) is preserved. This can be seen from

$$\begin{aligned}
\mathrm{tr}_1 L'_{i-1} &= \mathrm{tr}_1[(A_{i-2}^{-\top} \otimes A_{i-1})L_{i-1}(A_{i-2}^\top \otimes A_{i-1}^{-1})] \\
&= A_{i-1}(\mathrm{tr}_1 L_{i-1})A_{i-1}^{-1} \\
&= A_{i-1}(\mathrm{tr}_2 L_i)A_{i-1}^{-1} \\
&= \mathrm{tr}_2[(A_{i-1}^{-\top} \otimes A_i)L_i(A_{i-1}^\top \otimes A_i^{-1})] \\
&= \mathrm{tr}_2 L'_i.
\end{aligned} \tag{3.69}$$

On the other hand, if $n_i \neq s_{i-1}s_i$ for some i , then it is possible to show that $\Pi_{s_1 \dots s_{d-1}}$ is not reductive with the same argument as in proposition 3.4. \square

3.4.6 Geodesics

By an argument completely analogous to proposition 3.5, we have the following result.

Proposition 3.15. *Given a tangent vector $X \in T_p \Pi_{s_1 \dots s_{d-1}}$, $\exp_p(X)$ can be estimated using*

$$\begin{aligned}
&\frac{110}{3}n_1s_1^2 + (146 + 36z_1)s_1^3 + \mathcal{O}(n_1s_1 + s_1^2) \\
&+ \sum_{i=2}^{d-1} \frac{110}{3}n_i(s_{i-1}s_i)^2 + (146 + 36z_i)(s_{i-1}s_i)^3 + \mathcal{O}(n_i s_{i-1}s_i + (s_{i-1}s_i)^2) \\
&+ \frac{110}{3}n_d s_{d-1}^2 + (146 + 36z_d)s_{d-1}^3 + \mathcal{O}(n_d s_{d-1} + s_{d-1}^2) \text{ basic operations}
\end{aligned} \tag{3.70}$$

where $z_i = \lceil \log_2 \|X_i g_i^{-1}\| \rceil + 2$.

3.5 Application: low-rank tensor approximation

Recall from section 2.5.3 that we can get an approximate CP decomposition to a tensor from solving an optimization problem. Recall from (2.42) the definition of the objective function f , that T is a tensor in $\mathbb{R}^{n_1} \otimes \dots \otimes \mathbb{R}^{n_d}$, and the

notation $V_i = [v_i^1 \ \cdots \ v_i^r]$ for the factors. One immediate application of the Riemannian structure that we have derived is to implement *Riemannian gradient descent* (RGD) for f .

It is straightforward to show local convergence for such algorithms, in the same way as for classical gradient descent. Absil, Mahony, and Sepulchre [AMS08, chapter 4] and Boumal [Bou23, chapter 4] are standard resources on this topic.

The *gradient* is defined in terms of the Levi-Civita connection,

$$\text{grad } f(Y) = \nabla_Y f \quad \forall Y \in \mathcal{X}(M). \quad (3.71)$$

It is a cotangent vector. To get a direction on the manifold, we need to find the tangent vector X satisfying the condition $\langle X, \cdot \rangle = \text{grad } f$. In practice, we need find an expression to invert the metric. Let ξ be the coordinate expression for $\text{grad } f$. The condition then becomes

$$\text{tr} \left(X_i \begin{bmatrix} 1 & \Gamma_{12} \end{bmatrix} \begin{bmatrix} g_{11} & \\ g_{21} & 1 \end{bmatrix}^{-1} \begin{bmatrix} g_{11} & \\ g_{21} & 1 \end{bmatrix}^{-\text{T}} \begin{bmatrix} 1 & \Gamma_{12} \end{bmatrix}^{\text{T}} Y_i^{\text{T}} \right) = \text{tr}(\xi_i Y_i^{\text{T}}), \quad (3.72)$$

which simplifies to

$$\text{tr}(X_i [(1 - \Gamma_{12} g_{21}) g_{11}^{-1} g_{11}^{-\text{T}} (1 - \Gamma_{12} g_{21})^{\text{T}} + \Gamma_{12} \Gamma_{12}^{\text{T}}] Y_i^{\text{T}}) = \text{tr}(\xi_i Y_i^{\text{T}}). \quad (3.73)$$

We can make a substitution $Y_i \rightsquigarrow Y_i'$ to get

$$\text{tr}(X_i Y_i^{\text{T}}) = \text{tr}(\xi_i [(1 - \Gamma_{12} g_{21}) g_{11}^{-1} g_{11}^{-\text{T}} (1 - \Gamma_{12} g_{21})^{\text{T}} + \Gamma_{12} \Gamma_{12}^{\text{T}}]^{-1} Y_i'^{\text{T}}). \quad (3.74)$$

And so we find that

$$X_i = \xi_i [(1 - \Gamma_{12} g_{21}) g_{11}^{-1} g_{11}^{-\text{T}} (1 - \Gamma_{12} g_{21})^{\text{T}} + \Gamma_{12} \Gamma_{12}^{\text{T}}]^{-1}. \quad (3.75)$$

Given a gradient, we now know how to get the direction of steepest descent. Absil et al. [AMS08, section 3.6] discuss the correspondence between gradients and steepest descent directions more in depth in the context of matrix manifolds.

Moreover, the gradient of f is straightforward to compute using the product rule,

$$\begin{aligned} \nabla_Y f &= \nabla_Y \langle T - \llbracket V_1, \dots, V_d \rrbracket, T - \llbracket V_1, \dots, V_d \rrbracket \rangle \end{aligned} \quad (3.76)$$

$$= -2 \langle \llbracket Y_1, \dots, V_d \rrbracket, T - \llbracket V_1, \dots, V_d \rrbracket \rangle - \cdots - 2 \langle \llbracket V_1, \dots, Y_d \rrbracket, T - \llbracket V_1, \dots, V_d \rrbracket \rangle. \quad (3.77)$$

We implement our RGD with a *Armijo line search* to determine the step size. See for example Boumal [Bou23, algorithm 4.2]. Figure 3.1 shows a numerical experiment where we low-rank approximate the $10 \times 10 \times 10$ tensor

$$T$$

$$\begin{aligned} &= e_1 \otimes e_1 \otimes e_1 + e_2 \otimes e_4 \otimes e_6 + e_3 \otimes e_8 \otimes e_7 + e_4 \otimes e_3 \otimes e_8 + e_5 \otimes e_2 \otimes e_{10} \\ &+ e_6 \otimes e_6 \otimes e_2 + e_7 \otimes e_5 \otimes e_9 + e_8 \otimes e_7 \otimes e_3 + e_9 \otimes e_{10} \otimes e_4 + e_{10} \otimes e_9 \otimes e_5. \end{aligned} \tag{3.78}$$

The error is just the value of the objective function, f . We also compare to the standard ALS algorithm that we discussed in section 2.5.3, as well as to the theoretical best approximation. Note that T is exactly rank 10, $\text{mrnk}(10, 10, 10)$. For this specific problem, RGD consistently beats ALS but fails to converge to the best approximation for larger r .




Figure 3.1: Errors for different low-rank approximations of (3.78).

We also present a situation where RGD performs poorly. In figure 3.2, we approximate the $10 \times 10 \times 10$ tensor

$$T_{ijk} = \frac{1}{1 + (i/10)^2 + (j/10)^2 + (k/10)^2}. \tag{3.79}$$

For this problem, RGD clearly fails to converge for larger r . The error even grows with r .

We stress that while ALS is a very common approach to CP low-rank approximation, it is not necessarily state of the art, and that this subsection is meant as a proof-of-concept. Some more competitive Riemannian approaches are for example; the Gauss–Newton algorithm by Phan, Tichavsky, and Cichocki [PTC13]; the Gauss–Newton and Levenberg–Marquardt algorithms by




Figure 3.2: Errors for different low-rank approximations of (3.79).

Sorber, Van Barel, and De Lathauwer [SVD13]; and the trust region algorithm by Breiding and Vannieuwenhoven [BV18]. See also the survey of classical algorithms by Tomasi and Bro [TB06].

3.5.1 Degeneracy and regularization

While our naive RGD is not meant to be state of the art, it is still valuable to try to understand what goes wrong between figure 3.1 and figure 3.2. Since our RGD only uses first-order information, we are subject to the usual obstructions caused by saddle points and local minima. There are however other obstructions that can appear for this problem specifically, that have to do with conditioning.

In the tensor low-rank approximation literature, *degeneracy*, or *swamps*, refers to the observed phenomenon that terms tend to diverge at the same time, while almost cancelling. The phenomenon has been investigated in the context of chemometrics by Mitchell and Burdick [MB94], by Paatero [Paa00], and by Comon, Luciani, and De Almeida [CLD09]; and in the context of psychometrics by Stegeman [Ste06; Ste07; Ste08] and by Krijnen, Dijkstra, and Stegeman [KDS08]. The phenomenon is explained by Kruskal, Harshman, and Lundy [KHL89] and de Silva and Lim [SL08] as border rank effects.

Degeneracy is bad for conditioning, as rounding errors grow out of control [KB09]. We would like to avoid it if we can. One common way to avoid it is to restrict to nonnegative tensors. No catastrophic cancellations can occur when adding only positive numbers. This is however not always an option. For example, in our setting, it ruins the homogeneous structure. Another approach is to use a *regularizer*. For example, Navasca, De Lathauwer, and Kindermann [NDK08] introduce a Tikhonov-type regularizer, where an l^2

penalty term is added that depends on the previous iterate. Yet another option is *rank freezing*, introduced by Vermeylen and Vervliet, and De Lathauwer [VVD25], where a term is “frozen” temporarily if it seems to be converging.

3.6 Conclusion and outlook

In summary, we have solved for a homogeneous manifold structure of three families of low-rank manifolds: the set Σ_r of fixed-CP-rank tensors, the set $\Lambda_{t_1\dots t_d}$ of fixed-multilinear-rank tensors, and the set $\Pi_{s_1\dots s_d}$ of fixed-TT-rank tensors.

The conditions for when our construction works is essentially that the rank is low enough for Σ_r and $\Pi_{s_1\dots s_d}$, and for $\Lambda_{t_1\dots t_d}$ the condition is that one of the t 's is the product of the rest. This last case is much less permissive, but we argue that it is natural in applications where one of the modes corresponds to time.

The main motivation was to derive a Riemannian structure on these manifolds, where geodesics can be described explicitly. These geodesics are induced by geodesics on the general linear group, which can be expressed in terms of matrix exponentials. We also showed that we often don't have to do computations in the full general linear group, but rather that we can often do computations with much smaller matrices using the *vertical and horizontal* decomposition, a standard tool from Lie group theory. The main result that enables this is that there is an explicit rank decomposition of the horizontal space, something that happens if we carefully choose our representatives of points on the manifold.

Having solved for the geodesics on this manifold, there are several immediate applications in manifold optimization and statistics. We presented one such application, implementing a gradient descent method that descends along geodesics on the low-rank manifold to find low-rank approximations of a given tensor. The naive implementation left a lot to be desired, and we discussed some different reasons, like *degeneracy* and *swamps*, why this might be the case.

Another idea that we touched on briefly, but didn't explore here, is to define a low-rank retraction by taking a retraction on the action group and to let this induce a retraction on the quotient. Absil and Oseledets [AO15] published in 2015 a survey of 9 different low-rank retractions. A retraction based on the general linear group geodesics would add a new type of low-rank retraction.

Another natural application that we didn't explore, is to *Lie group integrators*. Having solved for the homogeneous structure of low-rank tensors,

we can now take any low-rank differential equation and lift it to the action group. Low-rank differential equations appear in *dynamical low-rank approximation*, a research area that is currently very active. This idea is the subject of ongoing research with researchers at NTNU: Elena Celledoni, Brynjulf Owren, and Jacob Goodman.

Recall from the start of the chapter that I contributed all parts of the chapter.

3.7 Appendix: error bound for Padé approximant

Recall the definition of ψ_1 from section 3.2.6 and the definition of r_{66} from (3.43), and let $\|\cdot\|$ be a matrix norm.

Lemma 3.16. *If $\|M\| \leq 1/2$, then $r_{66}(M)$ approximates $\psi_1(M)$ to within double precision $2^{-53} \approx 10^{-16}$.*

Proof. $r_{66}(M)$ has a Taylor expansion $\sum_{i=1}^{\infty} \frac{r_{66}^{(i)}(0)}{i!} M^i$. Since r_{66} is the Padé approximant to $\psi_1(M) = \sum_{i=1}^{\infty} \frac{1}{(i+1)!} M^i$, they agree up to the first 12 terms. We thus want to bound the series

$$\psi_1(M) - r_{66}(M) = \sum_{n=13}^{\infty} \frac{1 - (n+1)r_{66}^{(n)}(0)}{(n+1)!} M^n. \quad (3.80)$$

First, we compute term 13 and 14 manually. They are $M^{13}/149597947699200$ and $181M^{14}/29171599801344000$ respectively.

Second, we note that the n th derivative of r_{66} is a rational function $p_n(M)/q_n(M)$. The quotient rule gives the recurrence relation $p_{n+1} = p'_n q_n - p_n q'_n$ and $q_{n+1} = q_n^2$. We have that $q_n(0) = 1$ for all n . Moreover, $\|p_0^{(k)}(0)\|, \|q_0^{(k)}(0)\| \leq 1/2^k$ for all k . Using the triangle inequality in the

recurrence relation, this implies that $\|r_{66}^{(n)}(0)\| = \|p_n(0)\| \leq 1$. Thus

$$\begin{aligned}
\|\psi_1(M) - r_{66}(M)\| &= \left\| \sum_{n=13}^{\infty} \frac{1 - (n+1)r_{66}^{(n)}(0)}{(n+1)!} M^n \right\| \\
&\leq \frac{\|M\|^{13}}{149597947699200} + \frac{181\|M\|^{14}}{29171599801344000} \\
&\quad + \sum_{n=15}^{\infty} \frac{1 + (n+1)\|r_{66}^{(n)}(0)\|}{(n+1)!} \|M\|^n \\
&\leq \frac{(1/2)^{13}}{149597947699200} + \frac{181(1/2)^{14}}{29171599801344000} \\
&\quad + \sum_{n=15}^{\infty} \frac{1 + (n+1)}{(n+1)!} (1/2)^n \\
&= 3\sqrt{e} - \frac{2364006584786656554317}{477947491145220096000} \\
&\leq 2.7 \times 10^{-17}. \tag{3.81}
\end{aligned}$$

□

Lemma 3.17. *If $\|M\| \leq 1/2$, then $\text{mexp}(M)$ is approximated by its Padé approximant to within double precision.*

We do not prove lemma 3.17 here, but instead refer to Higham [Hig08, Section 10.3].

3.8 Appendix: counting the operations

We now prove proposition 3.5.

Proof. Recall the definition of A , B , A' , B' , and Γ_{12} from section 3.2.6. We then have

- $20n_i r^2/3 + 22r^3/3$ operations to build Γ_{12} using (3.24):
 - $8n_i r^2/3$ operations to divide g_{21} by g_{11} ,

- $2n_i r^2$ operations to multiply $g_{11}^{-\top} g_{21}^{\top}$ with $g_{21} g_{11}^{-1}$,
- $2r^3$ operations to square $g_{11}^{-\top} g_{21}^{\top} g_{21} g_{11}^{-1}$,
- $8r^3/3$ operations to divide by $1 + g_{11}^{-\top} g_{21}^{\top} g_{21} g_{11}^{-1}$,
- $8r^3/3$ operations to divide outside the parenthesis from the left by g_{11} ,
- $2n_i r^2$ to operations to multiply from the right by $g_{11}^{-\top} g_{21}^{\top}$,
- no extra operations to build A ,
- $2n_i r^2 + 8r^3/3$ operations to build B :
 - $2n_i r^2$ operations to multiply Γ_{12} with g_{21} ,
 - $8r^3/3$ operations to divide $1 - \Gamma_{12} g_{21}$ by g_{11} from the right,
- $2n_i r^2$ operations to multiply B and A ,
- $(52/3 + 4(z_i - 1))r^3$ operations to evaluate ψ_1 on BA using (3.43) and (3.44),
 - $12r^3$ operations to form the powers $2^{-z}M, \dots, (2^{-z}M)^6$,
 - $8r^3/3$ operations to evaluate the quotient,
 - $8r^3/3$ operations to evaluate $\text{mexp}(2^{-z}M)$,
 - $2(z - 1)r^3$ operations to form $\text{mexp}(2^{-z+1}M), \dots, \text{mexp}(2^{-1}M)$ by recursively using $\text{mexp}(2M) = \text{mexp}(M) \text{mexp}(M)$,
 - $2(z - 1)r^3$ operations to multiply the factors in (3.44),
- no extra operations to build A' or B' ,
- $2n_i(2r)^2$ operations to multiply B' and A' ,
- $(52/3 + 4(z_i - 1))(2r)^3$ operations to evaluate ψ_1 on $B'A'$,
- no extra operations to build the first term in (3.42),
- $4n_i r^2 + 2r^3$ operations to build the second term in (3.42):
 - $2n_i r^2$ operations to multiply B with $\begin{bmatrix} g_{11} \\ g_{21} \end{bmatrix}$,
 - $2r^3$ operations to multiply $\psi_1(BA)$ with $B \begin{bmatrix} g_{11} \\ g_{21} \end{bmatrix}$,
 - $2n_i r^2$ operations to multiply with A from the left,
- $8n_i r^2 + 8r^3$ operations to build the third term in (3.42):

- $2n_i(2r)r$ operations to multiply B' with $\begin{bmatrix} g_{11} \\ g_{21} \end{bmatrix}$,
- $2(2r)^2r$ operations to multiply $\psi_1(B'A')$ with $B' \begin{bmatrix} g_{11} \\ g_{21} \end{bmatrix}$,
- $2n_i(2r)r$ operations to multiply with A' from the left,
- $6n_i r^2 + 6r^3$ operations to build the last term in (3.42):
 - $2n_i r(2r)$ operations to multiply B with A' ,
 - $2r(2r)r$ operations to multiply BA' with $\psi_1(B'A')B' \begin{bmatrix} g_{11} \\ g_{21} \end{bmatrix}$,
 - $2r^3$ operations to multiply with $\psi_1(BA)$ from the left,
 - $2n_i r^2$ operations to multiply with A from the left.

□

4. Segre–Veronese manifolds

This chapter is based on arXiv preprint 2410.00664 “Warped geometries of Segre–Veronese manifolds” [Jac+24]. The preprint is in turn based on a chapter from Lars Swijsen’s PhD thesis [Swi22, chapter 6], where he computes the logarithmic map on the Segre manifold. The ideas were extended to more general warping factors and to the more general Segre–Veronese manifolds by Nick Vannieuwenhoven. My contribution to the paper was finalizing the logarithmic map on the pre-Segre–Veronese (theorem 4.7), the distance formula between compatible points (proposition 4.8) which in turn proves that distances are a monotone increasing function of the spherical distance (corollary 4.8.1) and that the pre-Segre–Veronese logarithmic map can be pushed forward to Segre–Veronese geodesics (theorem 4.11), the current version of the proof of distance between incompatible points (proposition 4.9), improving the matchmaking algorithm (proposition 4.14), classifying when the Segre–Veronese is geodesically connected (proposition 4.15), the sectional curvatures (section 4.5), and implementing the exponential and logarithmic maps in code; as well as performing the numerical experiments (section 4.6).

4.1 Introduction

The Segre–Veronese manifold is the set of partially symmetric rank-1 tensors. The *algebraic* geometry of the Segre–Veronese manifold and its secant varieties have been thoroughly studied since the 19th century [Ber+18]. By contrast, comparatively little is known about the *Riemannian* geometry of Segre–Veronese manifolds, let alone of (the smooth loci of) its higher secant varieties. To our knowledge, the Euclidean geometry of the Segre manifold as a submanifold of $\mathbb{R}^{n_1 \times \dots \times n_d}$ was only investigated in the prior works [SvV22; Swi22].

We study the Segre–Veronese manifold as a Riemannian manifold. We are primarily interested in its *geodesics* under a one-parameter family of Riemannian

metrics that includes the Euclidean metric. Our main contribution is analytic expressions for the exponential and logarithmic maps, in corollary 4.3.2 and theorem 4.11 respectively. We also establish a few auxiliary results, such as a neat formula for the distance function in proposition 4.13.

The manifold exponential was derived for the Segre manifold by Swijsen, Van der Veken, and Vannieuwenhoven [SvV22]. Their result generalizes fairly immediately to our more general setting, but our formula for the logarithmic map is novel. We also show that not every pair of rank-1 tensors on a Segre-Veronese manifold can be connected by a geodesic in the standard Euclidean geometry. We say that the manifold is not *geodesically connected*. On the other hand, some members of our one-parameter family are geodesically connected. We (partially) classify those in proposition 4.6. In some applications, such as computing *Karcher means*, it is more important that any pair of points can be connected by a geodesic than that Euclidean geometry is used.

We also implement the exponential and logarithmic maps in the Julia programming language and perform some numerical experiments in section 4.6. The implementation is available in the package `Manifolds.jl` [ABB24].

4.2 α -warped geometries

The Segre-Veronese manifold with the Euclidean metric is identified as a *warped product*. This identification then allows for a natural one-parameter generalization with different warping factors.

4.2.1 Warped geometries of \mathbb{R}_*^n

The punctured space \mathbb{R}_*^n is an open submanifold of \mathbb{R}^n . Hence, it is a smooth manifold. It can be equipped with the standard Euclidean metric, so that $\langle \dot{\mathbf{x}}, \dot{\mathbf{y}} \rangle_{\mathbf{a}} = \dot{\mathbf{x}} \cdot \dot{\mathbf{y}}$, where $\dot{\mathbf{x}}, \dot{\mathbf{y}} \in T_{\mathbf{a}}\mathbb{R}_*^n = \mathbb{R}^n$. Other meaningful geometries of \mathbb{R}_*^n exist. For example, \mathbb{R}_*^n is diffeomorphic to the product of the positive real line and an $(n-1)$ -dimensional sphere: $\mathbb{R}_*^n \simeq \mathbb{R}_{>0} \times S^{n-1}$. Under this identification, the above standard Euclidean metric can be expressed as

$$\langle \dot{x}\mathbf{u} + \lambda\dot{\mathbf{v}}, \dot{y}\mathbf{u} + \lambda\dot{\mathbf{w}} \rangle_{\lambda\mathbf{u}} = \dot{x}\dot{y} + \lambda^2\dot{\mathbf{v}} \cdot \dot{\mathbf{w}},$$

where $\dot{x}, \dot{y} \in \mathbb{R} = T_{\lambda}\mathbb{R}_{>0}$ and $\dot{\mathbf{v}}, \dot{\mathbf{w}} \in \mathbf{u}^\perp = T_{\mathbf{u}}S^{n-1}$. Note in particular that the metric that \mathbb{R}_*^n inherits from \mathbb{R}^n is *not* equal to the product metric, $\dot{x}\dot{y} + \dot{\mathbf{v}} \cdot \dot{\mathbf{w}}$, of $\mathbb{R}_{>0}$ and S^{n-1} . Instead, it is a so-called *warped product* [ONe83, Chapter 7] with the identity map as *warping function*.

Rather than considering only the Euclidean geometry of \mathbb{R}_*^n , we also consider its α -warped geometries; that is,

$$\langle \dot{x}\mathbf{u} + \lambda\dot{\mathbf{v}}, \dot{y}\mathbf{u} + \lambda\dot{\mathbf{w}} \rangle_{\lambda\mathbf{u}}^\alpha = \dot{x}\dot{y} + (\alpha\lambda)^2 \dot{\mathbf{v}} \cdot \dot{\mathbf{w}}. \tag{4.1}$$

The 1-warped geometry is the Euclidean geometry of \mathbb{R}_*^n , with straight lines as geodesics between points. If $\alpha > 1$, then distances in the spherical directions will be elongated, implying that geodesics between points will tend to be forced toward the origin. On the other hand, if $0 < \alpha < 1$, then spherical movements are easier, implying that geodesics between points will bend away from the origin. This is visualized in figure 4.1. The precise form of the geodesics is stated in proposition 4.3.




Figure 4.1: Geodesics between $(0, 1)$ and $(1, 0)$ in the α -warped geometry for $\alpha = 0.01$ (the red path approaching the semicircle), $\frac{1}{4}$, $\frac{1}{2}$, $\frac{3}{4}$, 1 (the purple straight line), $\frac{5}{4}$, $\frac{3}{2}$, $\frac{7}{4}$, and 1.99 (the blue path that narrowly avoids the puncture at the origin).

4.2.2 Warped geometries of the (pre-)Segre–Veronese manifold

\mathbb{R}_*^n is a trivial example of a Segre–Veronese manifold

$$\mathcal{S}_\pm^{\mathbf{k}} := \{ \lambda \mathbf{x}_1^{\otimes k_1} \otimes \cdots \otimes \mathbf{x}_d^{\otimes k_d} \mid \lambda \in \mathbb{R}_*, \mathbf{x}_i \in S^{n_i-1}, i = 1, \dots, d \}, \tag{4.2}$$

where $\mathbf{x}_i^{\otimes k_i}$ is the k_i -fold *tensor product* of \mathbf{x}_i with itself, i.e., $\mathbf{x}_i^{\otimes k_i} := \mathbf{x}_i \otimes \cdots \otimes \mathbf{x}_i$. Define moreover the “positive” Segre–Veronese manifold

$$\mathcal{S}^{\mathbf{k}} := \{ \lambda \mathbf{u}_1^{\otimes k_1} \otimes \cdots \otimes \mathbf{u}_d^{\otimes k_d} \mid \lambda \in \mathbb{R}_{>0}, \mathbf{u}_i \in S^{n_i-1}, i = 1, \dots, d \}, \quad (4.3)$$

If all k_i ’s are even, then $\mathcal{S}_{\pm}^{\mathbf{k}}$ consists of two connected components, the positive Segre–Veronese manifold $\mathcal{S}^{\mathbf{k}}$ and its mirror image along the origin, the “negative” Segre–Veronese manifold $-\mathcal{S}^{\mathbf{k}}$. This means that $p \in \mathcal{S}^{\mathbf{k}}$ and $q \in -\mathcal{S}^{\mathbf{k}}$ are not connected by any curve in $\mathcal{S}_{\pm}^{\mathbf{k}}$. In particular, there are no geodesics connecting them. Nevertheless, the two connected components are isometric. On the other hand, if at least one of the k_i ’s is odd, then it is easy to show that $\mathcal{S}_{\pm}^{\mathbf{k}} = \mathcal{S}^{\mathbf{k}} = -\mathcal{S}^{\mathbf{k}}$. That is, the “regular”, “positive”, and “negative” flavors of the Segre–Veronese manifold all coincide. Therefore, we will henceforth study the component $\mathcal{S}^{\mathbf{k}}$ without loss of generality.

Locally the Segre–Veronese manifold is diffeomorphic to $\mathbb{R}_{>0} \times (S^{n_1-1})^{\times k_1} \times \cdots \times (S^{n_d-1})^{\times k_d}$. We equip this with the α -warped metric, resulting in the α -warped *pre-Segre–Veronese*,

$$\mathcal{P}_{\alpha}^{\mathbf{k}} = \mathbb{R}_{>0} \times_{\alpha \text{ id}} (S^{n_1-1})^{\times k_1} \times \cdots \times (S^{n_d-1})^{\times k_d} \quad (4.4)$$

$$\langle (\dot{x}, \dot{\mathbf{u}}_1, \dots, \dot{\mathbf{u}}_d), (\dot{y}, \dot{\mathbf{v}}_1, \dots, \dot{\mathbf{v}}_d) \rangle_{(\lambda, \mathbf{u}_1, \dots, \mathbf{u}_d)} = \dot{x}\dot{y} + (\alpha\lambda)^2 \sum_{i=1}^d k_i \dot{\mathbf{u}}_i \cdot \dot{\mathbf{v}}_i. \quad (4.5)$$

Similarly to section 4.2.1, $\alpha = 1$ corresponds (locally) to the embedding in Euclidean space.

The main idea is that the tensor product takes the pre-Segre–Veronese manifold $\mathcal{P}_{\alpha}^{\mathbf{k}}$ in a nice way to the Segre–Veronese manifold $\mathcal{S}_{\alpha}^{\mathbf{k}}$. Specifically, it is a *normal Riemannian covering map* [ONe83], which means that (i) every point $\mathcal{T} \in \mathcal{S}_{\alpha}^{\mathbf{k}}$ has the same number, N , of preimages in $\mathcal{P}_{\alpha}^{\mathbf{k}}$ (ii) every point $\mathcal{T} \in \mathcal{S}_{\alpha}^{\mathbf{k}}$ has a neighborhood $U_{\mathcal{T}}$ such that its preimage $\otimes^{-1}(U_{\mathcal{T}})$ is a disjoint union of N domains, and (iii) on each of those domains, \otimes is an isometry. Property (iii) is called *local isometry*. A diffeomorphism $\iota: \mathcal{P}_{\alpha}^{\mathbf{k}} \rightarrow \mathcal{P}_{\alpha}^{\mathbf{k}}$ that satisfies $\otimes \circ \iota = \otimes$ is called a *deck transformation*. Informally, it shuffles the preimages of the covering map. A deck transform that is also an isometry is called an *isometric deck transform*.

Lemma 4.1. *The map*

$$\otimes: \mathcal{P}_{\alpha}^{\mathbf{k}} \rightarrow \mathcal{S}_{\alpha}^{\mathbf{k}}, \quad (\lambda, \mathbf{u}_1, \dots, \mathbf{u}_d) \mapsto \lambda \mathbf{u}_1^{\otimes k_1} \otimes \cdots \otimes \mathbf{u}_d^{\otimes k_d}$$

is a normal Riemannian covering where the isometric deck transforms are of the form $\iota_{\sigma}(\lambda, \mathbf{u}_1, \dots, \mathbf{u}_d) = (\lambda, \sigma_1 \mathbf{u}_1, \dots, \sigma_d \mathbf{u}_d)$ with $\sigma_i \in \{-1, 1\}$ and $\sigma_1^{k_1} \cdots \sigma_d^{k_d} = 1$.

See section 4.8 for a proof.

An elementary but crucial fact about normal Riemannian coverings concerns the lengths of curves in the covered manifold. The next result enables the computation of geodesics in $\mathcal{S}_\alpha^{\mathbf{k}}$ by studying the geodesics of the warped product $\mathcal{P}_\alpha^{\mathbf{k}}$.

Lemma 4.2. *Let $\phi: \mathcal{M} \rightarrow \mathcal{N}$ be a normal Riemannian covering. Let $\gamma: [0, 1] \rightarrow \mathcal{N}$ be a smooth curve. Let $p \in \mathcal{M}$ be any point such that $\phi(p) = \gamma(0)$. Then, there exists a unique smooth lift $\tilde{\gamma}: [0, 1] \rightarrow \mathcal{M}$ with $\tilde{\gamma}(0) = p$ and $\phi \circ \tilde{\gamma} = \gamma$ such that $\ell(\gamma) = \ell(\tilde{\gamma})$.*

In particular, a minimizing geodesic segment on \mathcal{N} lifts to a minimizing geodesic segment on \mathcal{M} .

See [ONe83, appendix A] for details.

By lemma 4.2, if we want to find a minimizing connecting geodesic between $P, Q \in \mathcal{S}_\alpha^{\mathbf{k}}$, it suffices to enumerate the connecting geodesics between the preimages and pick the shortest one. This is essentially what our proposed matchmaking algorithm will do in proposition 4.14.

4.3 The exponential map

Before presenting the main original contribution in section 4.4, we briefly investigate the exponential map on the (pre-)Segre–Veronese manifold. The following result is a straightforward generalization of [SvV22, Theorem 1] and its proof, which dealt only with the usual Euclidean geometry (i.e., $\alpha = 1$) of the Segre manifold (i.e., $\mathbf{k} = (1, \dots, 1)$). As we will rely on the general strategy of the proof, it is included for self-containedness.

Proposition 4.3 (Exponential map on $\mathcal{P}_\alpha^{\mathbf{k}}$). *Let*

$$p = (\lambda, \mathbf{u}_1, \dots, \mathbf{u}_d) \in \mathcal{P}_\alpha^{\mathbf{k}} \quad \text{and} \quad \dot{p} = (\dot{\lambda}, \dot{\mathbf{u}}_1, \dots, \dot{\mathbf{u}}_d) \in T_p \mathcal{P}_\alpha^{\mathbf{k}}.$$

Then

$$\exp_p(\dot{p}) = \left(\sqrt{(\lambda + \dot{\lambda})^2 + (\lambda\alpha N)^2}, \mathbf{u}_1 \cos a_1 + \frac{\dot{\mathbf{u}}_1}{\|\dot{\mathbf{u}}_1\|} \sin a_1, \dots, \mathbf{u}_d \cos a_d + \frac{\dot{\mathbf{u}}_d}{\|\dot{\mathbf{u}}_d\|} \sin a_d \right),$$

where

$$a_i = \frac{\|\dot{\mathbf{u}}_i\|}{\alpha N} \left(\frac{\pi}{2} - \arctan \left(\frac{\lambda + \dot{\lambda}}{\lambda \alpha N} \right) \right), \quad i = 1, \dots, d,$$

$$N = \sqrt{\sum_{i=1}^d k_i \|\dot{\mathbf{u}}_i\|^2}.$$

If $N = 0$ and $-\lambda < \dot{\lambda}$, the geodesics are straight lines:

$$\exp_p(\dot{p}) = (\lambda + \dot{\lambda}, \mathbf{u}_1, \dots, \mathbf{u}_d).$$

Proof. We are looking for the constant-speed geodesic γ with initial conditions $\gamma(0) = p$, $\gamma'(0) = \dot{p}$, so that we can evaluate it at 1.

If $N = 0$, the proof proceeds exactly as in [SvV22, Proof of Theorem 1]. So we only need to consider the case $N > 0$.

Let $\gamma(t) = (\lambda(t), \mathbf{u}_1(t), \dots, \mathbf{u}_d(t))$ be a geodesic in $\mathcal{P}_\alpha^{\mathbf{k}}$ with starting conditions $\gamma(0) = p$ and $\gamma'(0) = \dot{p}$. We apply the change of variable $t \mapsto s(t)$ with $s(0) = 0$ and $s'(0) > 0$ such that

$$\sigma(s) = (\mu(s), \mathbf{v}_1(s), \dots, \mathbf{v}_d(s)), \quad \text{with } \|(\mathbf{v}'_1(s), \dots, \mathbf{v}'_d(s))\|_{(\mathbf{v}_1(s), \dots, \mathbf{v}_d(s))} = 1, \quad (4.6)$$

i.e., σ is the geodesic satisfying $\sigma(s(t)) = \gamma(t)$ and which has unit *spherical* speed. Then, it holds that $\|\mathbf{v}'_i(0)\|_{\mathbf{v}_i(0)} s'(0) = \sqrt{k_i} \|\dot{\mathbf{u}}_i\|$, $i = 1, \dots, d$. From this we find that $s'(0) = N$ where N is as in the statement of the proposition. Then, [ONe83, Remark 7.39] implies that the projection of this geodesic on $S^{n_1-1} \times \dots \times S^{n_d-1}$ results in a geodesic. The geodesics on spheres are great circles:

$$\mathbf{v}_i(s) = \mathbf{u}_i \cos\left(\frac{\|\dot{\mathbf{u}}_i\|}{N} s\right) + \frac{\dot{\mathbf{u}}_i}{\|\dot{\mathbf{u}}_i\|} \sin\left(\frac{\|\dot{\mathbf{u}}_i\|}{N} s\right), \quad i = 1, \dots, d. \quad (4.7)$$

The coefficient $\mu: \mathbb{R} \rightarrow \mathbb{R}$ is determined by solving the *Euler-Lagrange equations* of the length functional, which in the α -warped geometry becomes

$$D = \int_0^1 \sqrt{\langle \sigma'(s), \sigma'(s) \rangle_{\sigma(s)}} ds = \int_0^1 \sqrt{\mu'(s)^2 + \alpha^2 \mu(s)^2} ds. \quad (4.8)$$

This ultimately leads to the differential equation $\alpha^2 \mu(s)^2 + 2\mu'(s)^2 - \mu(s)\mu''(s) = 0$ whose solution is of the form

$$\mu(s) = \frac{A}{\cos(\alpha s + B)}. \quad (4.9)$$

Solving for the constants A and B by exploiting the starting conditions, we arrive at $A = \lambda \cos(B)$ and $B = \arctan(\dot{\lambda}/(\alpha\lambda N))$. In this way, we obtain the geodesic given by (4.7) and (4.9).

The final step, consists of changing variables back to t . We have that

$$\begin{aligned} \|\gamma'(t)\| &= |s'(t)|\sqrt{\mu'(s(t))^2 + \alpha^2\mu(s(t))^2} \\ &= |s'(t)|\sqrt{\frac{\alpha^2 A^2 \sin^2(\alpha s(t) + B)}{\cos^4(\alpha s(t) + B)} + \frac{\alpha^2 A^2}{\cos^2(\alpha s(t) + B)}} \\ &= |s'(t)|\frac{\alpha A}{\cos^2(\alpha s(t) + B)}. \end{aligned} \quad (4.10)$$

Since $\sigma(s)$ is injective, the reparametrization $s(t)$ is a monotone increasing function. Especially, $s'(t) \geq 0$. By plugging $\|\gamma'(t)\| = \|\dot{p}\|_p = \sqrt{\dot{\lambda}^2 + (\alpha\lambda N)^2}$ into (4.10), we obtain the initial value problem

$$s(0) = 0, \quad s'(t) = \cos^2(\alpha s(t) + B) \frac{\sqrt{\dot{\lambda}^2 + (\alpha\lambda N)^2}}{\alpha A}. \quad (4.11)$$

Combining this with the initial condition $s(0) = 0$, we have a solution

$$s(t) = \alpha^{-1} \left(\frac{\pi}{2} - \arctan \left(\frac{\lambda + \dot{\lambda}t}{\alpha\lambda Nt} \right) \right). \quad (4.12)$$

Verifying that this solves (4.11) is a straightforward computation. For completeness, we include it in section 4.8.3.

To obtain the exponential map, we now just have to evaluate $\gamma(1)$. For the first component of $\gamma(1)$ we find, after some computations (see section 4.8.4), that

$$\mu(s(1)) = \sqrt{(\lambda + \dot{\lambda})^2 + (\alpha\lambda N)^2}. \quad (4.13)$$

The other components $\mathbf{v}_i(s(1))$ are obtained immediately by plugging the expression of $s(t)$ from (4.12) into (4.7). This concludes the proof. \square

From this result, it immediately follows that the exponential map is well defined almost everywhere.

Corollary 4.3.1. *The domain of the exponential map of $\mathcal{P}_\alpha^{\mathbf{k}}$ is the whole tangent space $T_p\mathcal{P}_\alpha^{\mathbf{k}}$, except for the half-line $\{(\dot{\lambda}, 0, \dots, 0) \mid \dot{\lambda} \leq -\lambda\}$.*

Having computed the geodesics on the pre-Segre–Veronese manifold, an application of lemmas 4.1 and 4.2 enables us to push them forward with \otimes to the Segre–Veronese manifold. Note that because Riemannian coverings are *local* isometries, we are only guaranteed that geodesics are mapped to geodesics, but not that minimizing geodesics are mapped to minimizing geodesics, i.e., the converse of lemma 4.2 is not true. In this way, we arrive at the geodesics and the following characterization of the exponential map of $\mathcal{S}_\alpha^{\mathbf{k}}$.

Corollary 4.3.2 (The exponential map on $\mathcal{S}_\alpha^{\mathbf{k}}$). *Let $\mathcal{P} \in \mathcal{S}_\alpha^{\mathbf{k}}$ be a rank-1 tensor and let $p = (\lambda, \mathbf{u}_1, \dots, \mathbf{u}_d) \in \otimes^{-1}(\mathcal{P})$ be any of its representatives on $\mathcal{P}_\alpha^{\mathbf{k}}$. Let $\hat{\mathcal{P}} \in T_p \mathcal{S}_\alpha^{\mathbf{k}}$, which can be expressed uniquely as*

$$\hat{\mathcal{P}} = \dot{\lambda} \mathcal{U} + \lambda \left(\nu_{k_1}(\dot{\mathbf{u}}_1) \otimes \mathbf{u}_2^{\otimes k_2} \cdots \otimes \mathbf{u}_d^{\otimes k_d} + \cdots + \mathbf{u}_1^{\otimes k_1} \otimes \cdots \otimes \mathbf{u}_{d-1}^{\otimes k_{d-1}} \otimes \nu_{k_d}(\dot{\mathbf{u}}_d) \right),$$

where

$$\mathcal{U} = \mathbf{u}_1^{\otimes k_1} \otimes \cdots \otimes \mathbf{u}_d^{\otimes k_d}, \quad (4.14)$$

$$\nu_{k_i}(\dot{\mathbf{u}}_i) = (\dot{\mathbf{u}}_i \otimes \mathbf{u}_i^{\otimes(k_i-1)}) + \cdots + (\mathbf{u}_i^{\otimes(k_i-1)} \otimes \dot{\mathbf{u}}_i) \quad \text{for all } \dot{\mathbf{u}}_i \in T_{\mathbf{u}_i} S^n. \quad (4.15)$$

Let $\hat{p} = (\dot{\lambda}, \dot{\mathbf{u}}_1, \dots, \dot{\mathbf{u}}_d) \in T_p \mathcal{P}_\alpha^{\mathbf{k}}$ be the corresponding lifted tangent vector. Then,

$$\exp_p(\hat{\mathcal{P}}) = \otimes(\exp_p(\hat{p})),$$

insofar as \hat{p} is in the domain of \exp_p .

Another interesting consequence of proposition 4.3 and corollary 4.3.1 is that it entails the next result about the radius of the largest open ball in $T_p \mathcal{S}_\alpha^{\mathbf{k}}$ for which the exponential map is a diffeomorphism, i.e., the local injectivity radius of the Segre–Veronese manifold.

Corollary 4.3.3 (Injectivity radius). *The local injectivity radius of $\mathcal{P}_\alpha^{\mathbf{k}}$, and thus also of $\mathcal{S}_\alpha^{\mathbf{k}}$, at $p = (\lambda, \mathbf{u}_1, \dots, \mathbf{u}_d)$ is λ . The (global) injectivity radius of $\mathcal{P}_\alpha^{\mathbf{k}}$ and $\mathcal{S}_\alpha^{\mathbf{k}}$, i.e., the infimum of the local injectivity radii over the manifold, is zero.*

Note that the latter claim is also immediate from the fact that neither $\mathcal{P}_\alpha^{\mathbf{k}}$ nor $\mathcal{S}_\alpha^{\mathbf{k}}$ are complete metric spaces. Indeed, the sequence $p_n = (n^{-1}, \mathbf{u}_1, \dots, \mathbf{u}_d) \in \mathcal{P}_\alpha^{\mathbf{k}}$ tends to $(0, \mathbf{u}_1, \dots, \mathbf{u}_d) \notin \mathcal{P}_\alpha^{\mathbf{k}}$ as $n \rightarrow \infty$, and the image of p_n in $\mathcal{S}_\alpha^{\mathbf{k}}$ tends to $0 \notin \mathcal{S}_\alpha^{\mathbf{k}}$.

Finally, we highlight that the exponential map on the pre-Segre–Veronese manifold $\mathcal{S}_\alpha^{\mathbf{k}}$ can be computed numerically very efficiently by straightforward implementation of the formulas in proposition 4.3. The cost is about $\mathcal{O}(\sum_{i=1}^d n_i)$ operations to compute N and all $\mathbf{u}_i(t)$'s, and a constant number of operations for $\lambda(t)$ and $g(t)$.

In the context of Riemannian optimization, the computational advantage and simplicity of using geodesics is clear with respect to the default alternative choice, applicable only to the standard Euclidean geometry (i.e., $\alpha = 1$), namely computing a quasi-optimal projection [Hac19, Chapter 10] with a truncated higher-order singular value decomposition [DDV00a], exploiting the structure of the tangent vectors, as performed in [KSV14].¹

4.4 The logarithmic map

4.4.1 Compatibility

One main motivation for considering α -warped geometries of the Segre–Veronese manifold is the observation in [Swi22, Chapter 6] that not all pairs of points on the Segre manifold are connected by a geodesic. That is, the Segre manifold is not geodesically connected in the standard Euclidean geometry.

The main result of this subsection, proposition 4.5, is that the geodesic connectedness of $\mathcal{P}_\alpha^{\mathbf{k}}$ can be characterized by the following property.

Definition 4.4 (α -compatibility). Consider two points $p = (\lambda, \mathbf{u}_1, \dots, \mathbf{u}_d)$ and $q = (\mu, \mathbf{v}_1, \dots, \mathbf{v}_d)$ in $\mathcal{P}_\alpha^{\mathbf{k}}$. Let

$$M(p, q) = \sqrt{\sum_{i=1}^d k_i \cdot \angle^2(\mathbf{u}_i, \mathbf{v}_i)}.$$

Then, p and q are called α -compatible if $M(p, q) < \alpha^{-1}\pi$.

The dependency of $M(p, q)$ on the parameters $\mathbf{k} \geq 1$ will not be emphasized in the notation, instead being implicit in the fact that $p, q \in \mathcal{P}_\alpha^{\mathbf{k}}$.

Proposition 4.5. *Two points p and q can be connected by a minimizing geodesic in $\mathcal{P}_\alpha^{\mathbf{k}}$ if and only if p and q are α -compatible.*

Proof. The case where $\mathbf{u}_i = \mathbf{v}_i$ for all $i = 1, \dots, d$ corresponds to the case of straight line geodesics (i.e., $M(p, q) = 0$) in proposition 4.3. All such p and

¹Since partially symmetric tensors form a subset of general tensors, a quasi-optimal projection to the Segre manifold that happens to result in a partially symmetric tensor will be quasi-optimal for the Segre–Veronese manifold as well. A rank- (r_1, \dots, r_d) (parallel) truncated higher-order singular value decomposition of a partially symmetric tensor will be partially symmetric, provided that the r_i th and $(r_i + 1)$ th singular values of the i th flattening (or matricization) are distinct.

q can be connected by a straight line and since $M(p, q) = 0$ they are always α -compatible. Hence, the claim holds in this special case.

The remainder of the proof treats the case $M(p, q) > 0$. Let $p = (\lambda, \mathbf{u}_1, \dots, \mathbf{u}_d) \in \mathcal{P}_\alpha^k$ and $q = (\mu, \mathbf{v}_1, \dots, \mathbf{v}_d) \in \mathcal{P}_\alpha^k$. For brevity, we let $M = M(p, q)$. We prove the two directions separately.

Geodesics \implies compatibility Assume that p and q are connected by a minimizing geodesic σ , parameterized to have unit spherical speed as in (4.6). Then, Chen's result [Che99, Lemma 4.4] implies that the projection of σ to the spherical part is a minimizing unit-speed geodesic. This geodesic is a product of arc segments of the great circles connecting \mathbf{u}_i and \mathbf{v}_i . The distance on $(S^{n_1-1})^{k_1} \times \dots \times (S^{n_d-1})^{k_d}$ between $(\mathbf{u}_1, \dots, \mathbf{u}_d)$ and $(\mathbf{v}_1, \dots, \mathbf{v}_d)$ thus satisfies $M^2 = \sum_{i=1}^d k_i \cdot \angle^2(\mathbf{u}_i, \mathbf{v}_i)$.

In section 4.2.1, we defined geodesics as critical points of a certain variational problem. Such a condition may equivalently be phrased in terms of an Euler-Lagrange equation such as (4.8). As σ is a geodesic, the proof of proposition 4.3 then implies that $\mu(s)$ is the form as in (4.9), where $A, B \in \mathbb{R}$ are determined by the conditions $\mu(0) = \lambda$ and $\mu(M) = \mu$. We have $A = \lambda \cos(B)$, as in the proof of proposition 4.3, and

$$\frac{\mu(0)}{\mu(M)} = \frac{\lambda}{\mu} = \frac{\cos(B + \alpha M)}{\cos(B)} = \cos(\alpha M) - \sin(\alpha M) \tan(B).$$

Hence,

$$B = \arctan\left(\cot(\alpha M) - \frac{\lambda\mu^{-1}}{\sin(\alpha M)}\right),$$

provided that $\alpha M \neq \pi k$ for $k \in \mathbb{Z}$. Since $\sigma(s)$ exists by assumption, the solution $\mu(s) > 0$ for all $0 \leq s \leq M$. The inverse of the cosine has singularities at $B + \alpha s = \frac{\pi}{2} + k\pi$ for $k \in \mathbb{Z}$, so $\mu(s)$ cannot pass through them and we must have $-\frac{\pi}{2} + k\pi < B + \alpha s < \frac{\pi}{2} + k\pi$ for all $0 \leq s \leq M$. In particular, the endpoint $s = 0$ implies that $-\frac{\pi}{2} + k\pi < B < \frac{\pi}{2} + k\pi$. Consequently, $0 < \alpha M < \pi$ because we also have that $M > 0$ and $\alpha > 0$. Hence, p and q are strictly α -compatible.

Compatibility \implies geodesics As before, consider minimizing geodesics $\mathbf{v}_i(s)$ connecting \mathbf{u}_i and \mathbf{v}_i on the sphere $(S^{n_i-1}, k_i \langle \cdot, \cdot \rangle)$ for $i = 1, \dots, d$. We can parameterize them so that they have unit spherical speed as in (4.6).

If there exists a geodesic connecting p and q , it must be a solution to the Euler-Lagrange equation. Proceeding as above, a solution $\mu(s)$ exists only if

$\alpha M \neq \pi k$ for $k \in \mathbb{Z}$. As p and q are α -compatible, $0 < \alpha M < \pi$, so $\mu(s) \subset \mathbb{R}$ is a valid smooth curve. For $\mu(s)$ to be a geodesic, however, it must also be contained in $\mathbb{R}_{>0}$ for all $0 \leq s \leq M$.

Since $-\frac{\pi}{2} < B < \frac{\pi}{2}$, it follows from $0 < \alpha M < \pi$ that

$$B + \alpha M = \arctan\left(\cot(\alpha M) - \frac{\lambda\mu^{-1}}{\sin(\alpha M)}\right) + \alpha M < \arctan(\cot(\alpha M)) + \alpha M,$$

because the arctangent is a monotonically increasing function and $\frac{\lambda\mu^{-1}}{\sin(\alpha M)} > 0$ when $\alpha M \in (0, \pi)$. Exploiting that $\cot(\alpha M) = \tan(\frac{\pi}{2} - \alpha M)$, we get $-\frac{\pi}{2} < B + \alpha M < \frac{\pi}{2}$ for all $0 < \alpha M < \pi$. Consequently, $\mu(s)$ is smooth curve that is strictly positive for all $0 \leq s \leq M$ and solves the Euler–Lagrange equation. Hence, $\sigma(s) = (\mu(s), \mathbf{v}_1(s), \dots, \mathbf{v}_d(s))$ is a geodesic connecting p and q . This proves the other direction and concludes the proof. \square

Proposition 4.5 implies that by changing the α parameter, we can control whether two points are geodesically connected. Specifically, we have the following result.

Proposition 4.6. *The pre-Segre–Veronese manifold $\mathcal{P}_\alpha^{\mathbf{k}}$ is geodesically connected if and only if $0 < \alpha < \frac{1}{\sqrt{k_1 + \dots + k_d}}$.*

Proof. Let $p, q \in \mathcal{P}_\alpha^{\mathbf{k}}$. Then, $M^2(p, q) \leq \sum_{i=1}^d k_i \pi^2$, with equality attained at $p = (\lambda, \mathbf{u}_1, \dots, \mathbf{u}_d)$ and $q = (\lambda, -\mathbf{u}_1, \dots, -\mathbf{u}_d)$. The result follows from proposition 4.5. \square

4.4.2 Pre-Segre–Veronese manifolds

Having determined when two points on the pre-Segre–Veronese manifold $\mathcal{P}_\alpha^{\mathbf{k}}$ can be connected by a minimizing geodesic, we proceed by computing such a geodesic between two α -compatible points.

Theorem 4.7 (Logarithmic map on $\mathcal{P}_\alpha^{\mathbf{k}}$). *Let*

$$p = (\lambda, \mathbf{u}_1, \dots, \mathbf{u}_d) \in \mathcal{P}_\alpha^{\mathbf{k}} \quad \text{and} \quad q = (\mu, \mathbf{v}_1, \dots, \mathbf{v}_d) \in \mathcal{P}_\alpha^{\mathbf{k}}$$

be such that $\mathbf{u}_i \neq -\mathbf{v}_i$ for all $i = 1, \dots, d$. Let $M = M(p, q)$ from definition 4.4. If p and q are α -compatible, then

$$\log_p(q) = (\mu \cos(\alpha M) - \lambda, \dot{\mathbf{u}}_1, \dots, \dot{\mathbf{u}}_d),$$

where

$$\dot{\mathbf{u}}_i := (\mathbf{v}_i - \langle \mathbf{u}_i, \mathbf{v}_i \rangle \mathbf{u}_i) \frac{\mu \sin(\alpha M) \angle(\mathbf{u}_i, \mathbf{v}_i)}{\lambda \alpha M \sin(\angle(\mathbf{u}_i, \mathbf{v}_i))}, \quad i = 1, \dots, d,$$

and the expressions $\sin(x)/x$ and $x/\sin(x)$ should be interpreted as 1 at $x = 0$. Especially, if $M = 0$, then $\log_p(q) = (\mu - \lambda, 0, \dots, 0)$.

Proof. We are looking for a geodesic γ with starting point p and ending point q .

The case $M = 0$ follows immediately from proposition 4.3. Hence, in the remainder of the proof we consider the case $M > 0$.

As in the proof of proposition 4.3, the idea is to use that the spherical part of γ is a geodesic on $S^{n_1-1} \times \dots \times S^{n_d-1}$. Since the logarithmic map on a sphere is known, we can pose as ansatz that $\log_p(q)$ is of the form

$$(\dot{\lambda}, \dot{\mathbf{u}}_1, \dots, \dot{\mathbf{u}}_d) = (C, D \log_{\mathbf{u}_1}(\mathbf{v}_1), \dots, D \log_{\mathbf{u}_d}(\mathbf{v}_d))$$

for some $C \in \mathbb{R}$ and $D > 0$. Now we just have to solve $q = \exp_p(\log_p q)$ for C and D .

Recall the expression for \exp_p from proposition 4.3. The first component yields

$$(\lambda + C)^2 + (\lambda \alpha D M)^2 = \mu^2, \quad (4.16)$$

having used that $\|\dot{\mathbf{u}}_i\| = D \|\log_{\mathbf{u}_i} \mathbf{v}_i\| = D \angle(\mathbf{u}_i, \mathbf{v}_i)$. Since $M > 0$, there exists at least one $1 \leq i \leq d$ so that $\mathbf{u}_i \neq \mathbf{v}_i$. In all of these coordinates, we get

$$\mathbf{v}_i = \mathbf{u}_i \cos a_i + \frac{\dot{\mathbf{u}}_i}{\|\dot{\mathbf{u}}_i\|} \sin a_i. \quad (4.17)$$

By projecting both sides to \mathbf{u}_i , we see that $\mathbf{u}_i \cdot \mathbf{v}_i = \cos(\angle(\mathbf{u}_i, \mathbf{v}_i)) = \cos a_i$, which implies $a_i = \angle(\mathbf{u}_i, \mathbf{v}_i)$. Then, from the expression of a_i in proposition 4.3, we obtain

$$\angle(\mathbf{u}_i, \mathbf{v}_i) = a_i = \frac{D \angle(\mathbf{u}_i, \mathbf{v}_i)}{\alpha D M} \left(\frac{\pi}{2} - \arctan \left(\frac{\lambda + C}{\alpha \lambda D M} \right) \right).$$

In this way we find

$$\cotan(\alpha M) = \frac{\lambda + C}{\alpha \lambda D M}. \quad (4.18)$$

The square system of polynomial equations (4.16) and (4.18) in C and D together with the constraint $D > 0$ defines the intersection of a circle with a line

on the half-plane $D > 0$. Ignoring the $D > 0$ constraint, the two intersection points are

$$C = \varsigma\mu \cos(\alpha M) - \lambda \quad \text{and} \quad D = \varsigma \frac{\mu \sin(\alpha M)}{\lambda \alpha M}, \quad \text{for } \varsigma \in \{-1, 1\}.$$

By α -compatibility we have $0 < \alpha M < \pi$, and since we posited $D > 0$, ς has to be 1.

With the above values for C and D , one verifies that $\exp_p \circ \log_p$ is the identity in the first component (4.16). The other components of $\exp_p \circ \log_p$ are also the identity, as can be verified by plugging the spherical logarithmic map [Bou23, Example 10.21]

$$\log_{\mathbf{u}_i}(\mathbf{v}_i) = (\mathbf{v}_i - (\mathbf{v}_i \cdot \mathbf{u}_i)\mathbf{u}_i) \frac{\sphericalangle(\mathbf{u}_i, \mathbf{v}_i)}{\sin(\sphericalangle(\mathbf{u}_i, \mathbf{v}_i))}$$

into (4.17). This concludes the proof. \square

Remark. We caution that if $\mathbf{u}_i = -\mathbf{v}_i$ for some i , then the logarithmic map is not well defined. However, p and q can still be connected by a geodesic. It is just that this geodesic is not unique. Choosing $\hat{\mathbf{u}}_i$ to be any vector orthogonal to \mathbf{u}_i and with norm π will produce such a geodesic.

An immediate consequence of the proof of the previous theorem is the following expression for the distance between points on the pre-Segre–Veronese manifold.

Proposition 4.8 (Distance between compatible points on $\mathcal{P}_\alpha^{\mathbf{k}}$). *Consider two points $p = (\lambda, \mathbf{u}_1, \dots, \mathbf{u}_d) \in \mathcal{P}_\alpha^{\mathbf{k}}$ and $q = (\mu, \mathbf{v}_1, \dots, \mathbf{v}_d) \in \mathcal{P}_\alpha^{\mathbf{k}}$ that are α -compatible. Let $M := M(p, q)$ from definition 4.4. Then,*

$$\text{dist}_{\mathcal{P}_\alpha^{\mathbf{k}}}(p, q) = \sqrt{\lambda^2 - 2\lambda\mu \cos(\alpha M) + \mu^2}. \tag{4.19}$$

Proof. For $M(p, q) = 0$, it follows immediately from proposition 4.3 that the distance between p and q on $\mathcal{P}_\alpha^{\mathbf{k}}$ is equal to $|\lambda - \mu|$. This coincides with (4.19).

Let $M(p, q) > 0$. The proof of theorem 4.7 shows that q is reached at $s = M(p, q)$ when the geodesic is parameterized to have unit speed on the spheres. Hence, we can use the reparameterization in (4.12) to determine the distance $\text{dist}_{\mathcal{P}_\alpha^{\mathbf{k}}}(p, q)$. It follows that it can be expressed as

$$\text{dist}_{\mathcal{P}_\alpha^{\mathbf{k}}}(p, q) = \lambda \cos(B)(\tan(B + \alpha M) - \tan B).$$

where $B = \arctan\left(\frac{\cos(\alpha M) - \lambda\mu^{-1}}{\sin(\alpha M)}\right)$. We proceed by simplifying this formula.

The addition formula for tan implies that

$$\begin{aligned} \text{dist}_{\mathcal{P}_\alpha^k}(p, q) &= \lambda \cos(B) \frac{\sin(B) \cos(\alpha M) + \cos(B) \sin(\alpha M)}{\cos(B) \cos(\alpha M) - \sin(B) \sin(\alpha M)} - \lambda \sin(B) \\ &= \frac{\lambda \sin(\alpha M)}{\cos(B) \cos(\alpha M) - \sin(B) \sin(\alpha M)}. \end{aligned}$$

We now substitute for B and use that $\cos(\arctan x) = 1/\sqrt{1+x^2}$ and $\sin(\arctan x) = x/\sqrt{1+x^2}$ to obtain

$$\begin{aligned} \text{dist}_{\mathcal{P}_\alpha^k}(p, q) &= \frac{\lambda \sin(\alpha M) \sqrt{1 + \left(\frac{\cos(\alpha M) - \lambda \mu^{-1}}{\sin(\alpha M)}\right)^2}}{\cos(\alpha M) - \left(\frac{\cos(\alpha M) - \lambda \mu^{-1}}{\sin(\alpha M)}\right) \sin(\alpha M)} \\ &= \frac{\lambda \sqrt{\sin^2(\alpha M) + (\cos(\alpha M) - \lambda \mu^{-1})^2}}{\lambda \mu^{-1}}, \end{aligned}$$

having used that $0 < \sin(\alpha M)$ because $0 < \alpha M < \pi$. The last equation is equivalent to (4.19), concluding the proof. \square

A formula, equivalent to (4.19), is

$$\text{dist}_{\mathcal{P}_\alpha^k}(p, q) = \sqrt{(\lambda - \mu)^2 + 4\lambda\mu \sin(\alpha M/2)}, \quad (4.20)$$

which is more numerically stable for points that are close together.

The distance between p and q on \mathcal{P}_α^k is thus given by the *law of cosines* applied to an imaginary triangle where one side has length λ , the other side has length μ , and the angle is the spherical distance $M(p, q)$ multiplied by the warping factor α . We thus immediately obtain the following result.

Corollary 4.8.1. *The distance $\text{dist}_{\mathcal{P}_\alpha^k}(p, q)$ is monotonically increasing as a function of the spherical distance $M(p, q)$ for $0 < \alpha M(p, q) < \pi$.*

Interestingly, we can also determine the distance between *incompatible* points, even though *minimizing* geodesics do not exist between them. While a minimizing geodesic does not exist, there exists a limiting piecewise smooth curve that is the infimizer of the length functional in (4.8). In figure 4.1, for example, this limiting piecewise smooth curve is the straight line from $(1, 0)$ to $(0, 0)$, composed with the straight line from $(0, 0)$ to $(0, 1)$.

Proposition 4.9 (Distance between incompatible points on $\mathcal{P}_\alpha^{\mathbf{k}}$). *Two α -incompatible points $p = (\lambda, \mathbf{u}_1, \dots, \mathbf{u}_d) \in \mathcal{P}_\alpha^{\mathbf{k}}$ and $q = (\mu, \mathbf{v}_1, \dots, \mathbf{v}_d) \in \mathcal{P}_\alpha^{\mathbf{k}}$ are at*

$$\text{dist}_{\mathcal{P}_\alpha^{\mathbf{k}}}(p, q) = \lambda + \mu.$$

Proof. Let γ_n be a sequence of unit speed curves in $\mathcal{P}_\alpha^{\mathbf{k}}$ that connect two incompatible points p and q , and whose lengths approach the infimum.

To lower bound $\ell(\gamma_n)$, note that γ_n must get arbitrarily close to 0 as $n \rightarrow \infty$, since otherwise the sequence lives in a compact subset (norm bounded from above and below, which is also a closed condition) of $\mathcal{P}_\alpha^{\mathbf{k}}$ and would thus have a (uniformly²) convergent subsequence. The limit of this subsequence would realize the infimum, and thus be a geodesic, contradicting our assumption about p and q being incompatible.

So fix n and let $a = (\epsilon, \mathbf{a}_1, \dots, \mathbf{a}_d) \in \mathcal{P}_\alpha^{\mathbf{k}}$ be a point on γ_n and let $\gamma^{(1)}$ denote the part of γ_n connecting p and a and let $\gamma^{(2)}$ be the part of γ_n connecting a and q . We have that

$$\ell(\gamma_n) = \ell(\gamma^{(1)}) + \ell(\gamma^{(2)}).$$

Moreover, if $\gamma^{(1)}(s) = (\nu(s), \mathbf{x}_1(s), \dots, \mathbf{x}_d(s))$, then

$$\ell(\gamma^{(1)}) = \int \sqrt{\nu'(s)^2 + \alpha^2 \nu(s)^2} ds \geq \int \sqrt{\nu'(s)^2} ds = \lambda - \epsilon$$

A similar argument shows that $\ell(\gamma^{(2)}) \geq \mu - \epsilon$. Especially, for every ϵ , there exists a n such that $\ell(\gamma_n) \geq \lambda + \mu - 2\epsilon$.

To upper bound $\inf_\gamma \ell(\gamma)$, consider the curve that goes first in a straight line from $p = (\lambda, \mathbf{u}_1, \dots, \mathbf{u}_d)$ to $(\epsilon, \mathbf{u}_1, \dots, \mathbf{u}_d)$, then in a circular arc from $(\epsilon, \mathbf{u}_1, \dots, \mathbf{u}_d)$ to $(\epsilon, \mathbf{v}_1, \dots, \mathbf{v}_d)$, and then in a straight line from $(\epsilon, \mathbf{u}_1, \dots, \mathbf{u}_d)$ to $(\epsilon, \mathbf{v}_1, \dots, \mathbf{v}_d)$ to $q = (\mu, \mathbf{u}_1, \dots, \mathbf{u}_d)$ to $(\epsilon, \mathbf{v}_1, \dots, \mathbf{v}_d)$. It has length $\lambda + \mu + (\alpha M - 2)\epsilon$.

Combining these two bounds and letting $\epsilon \rightarrow 0$ proves the proposition. \square

4.4.3 Segre–Veronese manifolds

By leveraging lemma 4.2 and theorem 4.7, minimizing geodesics on the Segre–Veronese manifold $\mathcal{S}_\alpha^{\mathbf{k}}$ that connect two rank-1 tensors are readily obtained.

²Unit speed implies equi-Lipschitz continuous, implies uniform convergence to a Lipschitz continuous curve.

Indeed, a minimizing geodesic connecting $\mathcal{P}, Q \in \mathcal{S}_\alpha^{\mathbf{k}}$ is the pushforward of any geodesic whose length coincides with

$$\min_{p \in \otimes^{-1}(\mathcal{P}), q \in \otimes^{-1}(Q)} \text{dist}_{\mathcal{P}_\alpha^{\mathbf{k}}}(p, q).$$

Corollary 4.8.1 then leads us naturally to the following terminology.

Definition 4.10 (Matched representatives). Let $\mathcal{P}, Q \in \mathcal{S}_\alpha^{\mathbf{k}}$. We say that the representatives $p^* \in \otimes^{-1}(\mathcal{P})$ and $q^* \in \otimes^{-1}(Q)$ are *matched* if

$$(p^*, q^*) \in \arg \min_{p \in \otimes^{-1}(\mathcal{P}), q \in \otimes^{-1}(Q)} M(p, q),$$

where M is the distance from definition 4.4.

We are now ready to state and prove the main result about minimizing geodesics connecting two rank-1 tensors on the α -warped Segre–Veronese manifold.

Theorem 4.11 (The logarithmic map of $\mathcal{S}_\alpha^{\mathbf{k}}$). *Let $\mathcal{P}, Q \in \mathcal{S}_\alpha^{\mathbf{k}}$ be rank-1 tensors. If $p = (\lambda, \mathbf{u}_1, \dots, \mathbf{u}_d) \in \otimes^{-1}(\mathcal{P})$ and $q = (\mu, \mathbf{v}_1, \dots, \mathbf{v}_d) \in \otimes^{-1}(Q)$ are matched, α -compatible representatives, then*

$$\log_p Q = (d_p \otimes)(\log_p q).$$

Proof. Since p and q are compatible, there is by definition 4.10 a minimizing geodesic between them. To pull this back to a minimizing geodesic on $\mathcal{S}_\alpha^{\mathbf{k}}$, we need to know that inequalities between the lengths are not reversed. That is, if γ and $\tilde{\gamma}$ are geodesics on $\mathcal{P}_\alpha^{\mathbf{k}}$ such that $(\otimes \circ \gamma)(0) = (\otimes \circ \tilde{\gamma})(0) = \mathcal{P}$, $(\otimes \circ \gamma)(1) = (\otimes \circ \tilde{\gamma})(1) = Q$, and $\ell(\gamma) \leq \ell(\tilde{\gamma})$, then we need to show that $\ell(\otimes \circ \gamma) \leq \ell(\otimes \circ \tilde{\gamma})$. This implication indeed holds as a direct consequence of corollary 4.8.1. \square

The following two results follow similarly.

Proposition 4.12. *Two rank-1 tensors \mathcal{P} and $Q \in \mathcal{S}_\alpha^{\mathbf{k}}$ can be connected by a minimizing geodesic in $\mathcal{S}_\alpha^{\mathbf{k}}$ if and only if \mathcal{P} and Q have matched α -compatible representatives.*

Proposition 4.13 (Distance between points on $\mathcal{S}_\alpha^{\mathbf{k}}$). *Let $\mathcal{P}, Q \in \mathcal{S}_\alpha^{\mathbf{k}}$. If $p = (\lambda, \mathbf{u}_1, \dots, \mathbf{u}_d) \in \otimes^{-1}(\mathcal{P})$ and $q = (\mu, \mathbf{v}_1, \dots, \mathbf{v}_d) \in \otimes^{-1}(Q)$ are matched, α -compatible representatives, then*

$$\text{dist}_{\mathcal{S}_\alpha^{\mathbf{k}}}(\mathcal{P}, Q) = \text{dist}_{\mathcal{P}_\alpha^{\mathbf{k}}}(p, q).$$

The straightforward approach to determine matched representatives of \mathcal{P} , $Q \in \mathcal{S}_\alpha^{\mathbf{k}}$ consists of choosing an arbitrary $p^* \in \otimes^{-1}(\mathcal{P})$ and then computing all distances $M(p^*, q)$ for all $q \in \otimes^{-1}(Q)$ and selecting a minimum. By lemma 4.1, the possible q can be labeled by $(\sigma_1, \dots, \sigma_d) \in \{-1, 1\}^d$ such that $\sigma_1^{k_1} \dots \sigma_d^{k_d} = 1$. Thus this strategy has a computational complexity of $\mathcal{O}(2^d \dim \mathcal{S}_\alpha^{\mathbf{k}})$, where $\mathcal{O}(\dim \mathcal{S}_\alpha^{\mathbf{k}})$ is the asymptotic cost for computing one distance $M(p^*, q)$. We now explain how a matched representative can be constructed efficiently in $\mathcal{O}(d \dim \mathcal{S}_\alpha^{\mathbf{k}})$ operations.

Let $q_{\text{ref}} \in \otimes^{-1}(Q)$ be a reference preimage. We then observe that

$$M(p^*, q)^2 = M(p^*, q_{\text{ref}})^2 + \sum_{i=1}^d \frac{1 - \sigma_i}{2} \Delta_i.$$

with $\Delta_i = k_i((\angle_i - \pi)^2 - \angle_i^2)$, where \angle_i is the angle between the i th components of p^* and q_{ref} . Minimizing $M(p^*, q)$ thus corresponds to maximizing $\sum_{i=1}^d \sigma_i \Delta_i$ over $\sigma_i \in \{-1, 1\}$. An unconstrained optimum is

$$\sigma_i = \begin{cases} 1, & \text{if } \Delta_i \text{ is nonnegative,} \\ -1, & \text{otherwise.} \end{cases} \quad (4.21)$$

If this assignment also satisfies the constraint $\sigma_1^{k_1} \dots \sigma_d^{k_d} = 1$, so that $q \in \otimes^{-1}(q_{\text{ref}})$, then we are done. Otherwise, there is at least one smallest $|\Delta_\ell|$ among the odd k_ℓ 's. The sign of the corresponding σ_ℓ should be flipped.

Proposition 4.14 (Matchmaking). *The algorithm described in the previous paragraph produces matched representatives.*

Proof. A detailed proof is given in section 4.8.2. □

The construction of the matchmaking algorithm suggests that determining an analogue of proposition 4.6 is more complicated for the Segre–Veronese manifold. Essentially, we have to consider the following problem: Given $p^* \in \mathcal{P}_\alpha^{\mathbf{k}}$, what is the set of $q \in \mathcal{P}_\alpha^{\mathbf{k}}$ such that the corresponding matched representatives $q^* \in \mathcal{P}_\alpha^{\mathbf{k}}$ are α -compatible with p^* , and when is this set the whole of $\mathcal{P}_\alpha^{\mathbf{k}}$? By our calculations, this question can be reduced to a combinatorial optimization problem. As we were unable to produce a closed form of its solution, we present only the following general result.

Proposition 4.15. *The Segre–Veronese manifold $\mathcal{S}_\alpha^{\mathbf{k}}$ is geodesically connected if $0 < \alpha < 1/\sqrt{k_1 + \dots + k_d}$ and is not geodesically connected if $2/\sqrt{k_1 + \dots + k_d} \leq \alpha$.*

Proof. The first statement is a corollary of lemma 4.1 and proposition 4.6.

To prove the second statement, assume $\alpha \geq 2/\sqrt{k_1 + \cdots + k_d}$ and pick $p, q \in \mathcal{P}_\alpha^k$ with angles $\sphericalangle_1 = \cdots = \sphericalangle_d = \frac{\pi}{2}$ between their components. These two representatives are matched. However, using $M(p, q)^2 = (k_1 + \cdots + k_d) \frac{\pi^2}{4}$ in definition 4.4 shows that they are not α -compatible. \square

For $1/\sqrt{k_1 + \cdots + k_d} < \alpha < 2/\sqrt{k_1 + \cdots + k_d}$, geodesic connectedness is more complicated and depends also on the parity of the k_i 's. The full analysis is beyond the scope of this article.

4.5 Sectional curvature

Formulae for the sectional curvature of a manifold are useful in many situations, for example when approximating or compressing data that lives on said manifold [Jac+25; DCN25; Zim20], or to ensure uniqueness of the Fréchet mean [Kar77; Ken90; Afs11], as required in the experiment from section 4.6. We compute these curvatures here for the α -warped Segre–Veronese manifolds.

Consider first the α -warped Segre manifold $\mathcal{S}_\alpha = \mathcal{S}_\alpha^{(1, \dots, 1)}$. To compute its curvature, we recall the following general result about warped products.

Proposition 4.16 (O’Neill [ONe83, p. 7.42]). *For $(a, b) \in \mathcal{M} \times_f \mathcal{N}$, let $\dot{\mathbf{x}}, \dot{\mathbf{y}}, \dot{\mathbf{z}} \in T_a \mathcal{M}$ and $\dot{\mathbf{u}}, \dot{\mathbf{v}}, \dot{\mathbf{w}} \in T_b \mathcal{N}$ be tangent vectors, which naturally lift to $T_{(a,b)}(\mathcal{M} \times_f \mathcal{N})$. Let $R_{\mathcal{M}}$ and $R_{\mathcal{N}}$ be the Riemannian curvature tensors on \mathcal{M} and \mathcal{N} , respectively. Then, the Riemann tensor R on $\mathcal{M} \times_f \mathcal{N}$ is described by:³*

1. $R(\dot{\mathbf{x}}, \dot{\mathbf{y}})\dot{\mathbf{z}} = R_{\mathcal{M}}(\dot{\mathbf{x}}, \dot{\mathbf{y}})\dot{\mathbf{z}}$,
2. $R(\dot{\mathbf{u}}, \dot{\mathbf{x}})\dot{\mathbf{y}} = -\frac{1}{f} \text{Hess}[f](\dot{\mathbf{x}}, \dot{\mathbf{y}})\dot{\mathbf{u}}$,
3. $R(\dot{\mathbf{x}}, \dot{\mathbf{y}})\dot{\mathbf{u}} = 0$,
4. $R(\dot{\mathbf{u}}, \dot{\mathbf{v}})\dot{\mathbf{x}} = 0$,
5. $R(\dot{\mathbf{x}}, \dot{\mathbf{u}})\dot{\mathbf{v}} = -\frac{\langle \dot{\mathbf{u}}, \dot{\mathbf{v}} \rangle}{f} \nabla_{\dot{\mathbf{x}}}(\text{grad } f)$,
6. $R(\dot{\mathbf{u}}, \dot{\mathbf{v}})\dot{\mathbf{w}} = R_{\mathcal{N}}(\dot{\mathbf{u}}, \dot{\mathbf{v}})\dot{\mathbf{w}} + \frac{\|\text{grad } f\|^2}{f^2} (\langle \dot{\mathbf{u}}, \dot{\mathbf{w}} \rangle \dot{\mathbf{v}} - \langle \dot{\mathbf{v}}, \dot{\mathbf{w}} \rangle \dot{\mathbf{u}})$.

³Note that we use a definition of R that is consistent with [Lee18] rather than [ONe83]. The result is that some signs differ in our expressions from theirs.

Herein, $\langle \cdot, \cdot \rangle$ is the inner product on $T_{(a,b)}(\mathcal{M} \times_f \mathcal{N})$, $\text{grad } f$ is the Riemannian gradient of f , $\text{Hess}[f]$ is the Riemannian Hessian of f , and ∇ is the Levi-Civita connection.

For \mathcal{S}_α , which locally looks like the warped product manifold \mathcal{P}_α , these expressions simplify considerably. Applying proposition 4.16 to $\mathcal{N} = S^{\mathbf{n},1} = S^{n_1-1} \times \cdots \times S^{n_d-1}$ and $\mathcal{M} = \mathbb{R}_{>0}$ while observing that $\text{grad } f = \alpha$ and $R_{\mathcal{M}} = 0$ yields the next result.

Corollary 4.16.1. *In the notation of proposition 4.16, the Riemann tensor R of \mathcal{S}_α at $\lambda \mathbf{u}_1 \otimes \cdots \otimes \mathbf{u}_d$ satisfies*

$$R(\dot{\mathbf{u}}, \dot{\mathbf{v}})\dot{\mathbf{w}} = R_{S^{\mathbf{n},1}}(\dot{\mathbf{u}}, \dot{\mathbf{v}})\dot{\mathbf{w}} + \frac{1}{\lambda^2}(\langle \dot{\mathbf{u}}, \dot{\mathbf{w}} \rangle \dot{\mathbf{v}} - \langle \dot{\mathbf{v}}, \dot{\mathbf{w}} \rangle \dot{\mathbf{u}}),$$

and it is 0 in the remaining directions.

As sectional curvatures can be defined in terms of the Riemann tensor, we quickly find the next consequence of the last result.

Corollary 4.16.2. *For $(\mathbf{u}_1, \dots, \mathbf{u}_d) \in S^{\mathbf{n},1}$, let the lifts of the tangent vectors $\dot{\mathbf{u}}_i \in T_{\mathbf{u}_i} S^{n_i-1}$ and $\dot{\mathbf{v}}_j \in T_{\mathbf{u}_j} S^{n_j-1}$ be orthonormal with respect to $\langle \cdot, \cdot \rangle_{\mathcal{P}_\alpha}$. The sectional curvature K of \mathcal{S}_α at $\lambda \mathbf{u}_1 \otimes \cdots \otimes \mathbf{u}_d$ satisfies*

$$K(\dot{\mathbf{u}}_i, \dot{\mathbf{v}}_j) = \begin{cases} \frac{1-\alpha^2}{\alpha^2 \lambda^2}, & \text{if } i = j, \\ -\frac{1}{\lambda^2}, & \text{otherwise,} \end{cases}$$

and $K = 0$ in all remaining directions.

Proof. Recall that $\langle \dot{\mathbf{u}}_i, \dot{\mathbf{v}}_j \rangle_{\mathcal{P}_\alpha} = \alpha^2 \lambda^2 \langle \dot{\mathbf{u}}_i, \dot{\mathbf{v}}_j \rangle_{S^{\mathbf{n},1}}$. We compute

$$\begin{aligned} K(\dot{\mathbf{u}}_i, \dot{\mathbf{v}}_j) &= \langle R(\dot{\mathbf{u}}_i, \dot{\mathbf{v}}_j)\dot{\mathbf{v}}_j, \dot{\mathbf{u}}_i \rangle_{\mathcal{P}_\alpha} \\ &= \langle R_{S^{\mathbf{n},1}}(\dot{\mathbf{u}}_i, \dot{\mathbf{v}}_j)\dot{\mathbf{v}}_j, \dot{\mathbf{u}}_i \rangle_{\mathcal{P}_\alpha} + \frac{1}{\lambda^2} \left(\langle \dot{\mathbf{u}}_i, \dot{\mathbf{v}}_j \rangle_{\mathcal{P}_\alpha}^2 - \langle \dot{\mathbf{v}}_j, \dot{\mathbf{v}}_j \rangle_{\mathcal{P}_\alpha} \langle \dot{\mathbf{u}}_i, \dot{\mathbf{u}}_i \rangle_{\mathcal{P}_\alpha} \right) \\ &= \alpha^2 \lambda^2 \langle R_{S^{\mathbf{n},1}}(\dot{\mathbf{u}}_i, \dot{\mathbf{v}}_j)\dot{\mathbf{v}}_j, \dot{\mathbf{u}}_i \rangle_{S^{\mathbf{n},1}} - \frac{1}{\lambda^2} \\ &= \frac{\alpha^2 \lambda^2}{\alpha^4 \lambda^4} \langle R_{S^{\mathbf{n},1}}(\alpha \lambda \dot{\mathbf{u}}_i, \alpha \lambda \dot{\mathbf{v}}_j) \alpha \lambda \dot{\mathbf{v}}_j, \alpha \lambda \dot{\mathbf{u}}_i \rangle_{S^{\mathbf{n},1}} - \frac{1}{\lambda^2} \\ &= \frac{1}{\alpha^2 \lambda^2} K_{S^{\mathbf{n},1}}(\alpha \lambda \dot{\mathbf{u}}_i, \alpha \lambda \dot{\mathbf{v}}_j) - \frac{1}{\lambda^2} \\ &= \frac{1}{\alpha^2 \lambda^2} \delta_{ij} - \frac{1}{\lambda^2}. \end{aligned}$$

In all other directions the sectional curvature is 0, because $R = 0$ in these directions. \square

The Segre–Veronese manifold $\mathcal{S}_\alpha^{\mathbf{k}}$, $\mathbf{k} = (k_1, \dots, k_d)$, is a Riemannian submanifold of the Segre manifold $\mathcal{S}_\alpha^{\mathbf{1}}$, where $\mathbf{1}$ is a vector of $k_1 + \dots + k_d$ ones. If we look closely at the expression in proposition 4.3 for geodesics on $\mathcal{S}_\alpha^{\mathbf{1}}$, we see that starting out in a tangent direction to $\mathcal{S}_\alpha^{\mathbf{k}} \subset \mathcal{S}_\alpha^{\mathbf{1}}$ produces a path wholly contained in $\mathcal{S}_\alpha^{\mathbf{k}}$. We have thus proved the following result.

Proposition 4.17. *The Segre–Veronese manifold $\mathcal{S}_\alpha^{\mathbf{k}}$ is a totally geodesic submanifold of the Segre manifold $\mathcal{S}_\alpha^{\mathbf{1}}$, and so its sectional curvature is described by restricting K in corollary 4.16.2 to the appropriate tangent subspace.*

4.6 Application: consensus aggregation

The α -warped Segre–Veronese manifold $\mathcal{S}_\alpha^{\mathbf{k}}$ was implemented in Julia, including connecting and starting geodesics and the geodesic distance. Our implementation is available as part of the Manifolds.jl library [ABB24].

As a numerical example illustration of the connecting geodesics in the α -warped geometry of the Segre manifold, we recall the *consensus aggregation problem* from Chafamo, Shanmugam, and Tokcan [CST23]: Given M approximate *tensor rank decompositions* of the same tensor $\mathcal{A} \in \mathbb{R}^{n_1 \times \dots \times n_d}$,

$$\mathcal{A} \approx \mathcal{A}^{(m)} := \sum_{i=1}^r \mathcal{A}_i^{(m)} := \sum_{i=1}^r \mathbf{a}_{1i}^{(m)} \otimes \dots \otimes \mathbf{a}_{di}^{(m)}, \quad m = 1, \dots, M,$$

estimate the true decomposition of \mathcal{A} .

To solve the above problem, a consensus aggregation algorithm was proposed in [CST23] that takes elementwise medians of corresponding vectors $\mathbf{a}_{li}^{(m)}$, $m = 1, \dots, M$, of matching summands in the decomposition. While tensor rank decompositions usually have generically unique decompositions into a set of rank-1 tensors [MM24; TC24; Bal24], to perform aggregation one needs to determine how these rank-1 tensors match up across multiple decompositions. For this reason, the rank-1 tensors are first matched using a (k -means) clustering algorithm in [CST23]. This strategy can work well if the decomposition problem is well-conditioned [BV18], so that $\mathcal{A}^{(m)}$ being close to \mathcal{A} will imply that the corresponding rank-1 tensors are close too. After matching up the rank-1 tensors, the results are aggregated in [CST23] by setting

$$\widehat{\mathbf{a}}_i^l = \text{median}(\mathbf{a}_{li}^{(m)} : m = 1, \dots, M), \quad i = 1, \dots, r, \quad l = 1, \dots, d.$$

The median is taken elementwise. Note that it is possible for $\widehat{\mathbf{a}}_i^l$ to have different values if different representatives are chosen. This is circumvented in [CST23] by aggregating positive data, for which there is a natural, positive, representative. The aggregated tensor decomposition is then

$$\mathcal{A} \approx \widehat{\mathcal{A}} := \sum_{i=1}^r \widehat{\mathbf{a}}_i^1 \otimes \cdots \otimes \widehat{\mathbf{a}}_i^d.$$

The above is a simplified description of the essential ingredients of [CST23, Section 3].

From a geometric viewpoint, it is natural to estimate each rank-1 tensor with the Fréchet mean on the warped Segre manifold \mathcal{S}_α . Recall that the Fréchet mean or Riemannian center of mass of points x_1, \dots, x_k on a Riemannian manifold \mathcal{M} is

$$\text{mean}_{\mathcal{M}}(x_1, \dots, x_k) = \arg \min_{p \in \mathcal{M}} \sum_{i=1}^n \text{dist}_{\mathcal{M}}(p, p_i)^2.$$

Its uniqueness has been investigated among others in [Kar77; Ken90; Afs11]. For manifolds with negative sectional curvatures, the Fréchet mean is unique and the above defines a valid function. Hence, we propose to estimate \mathcal{A} instead as

$$\mathcal{A} \approx \widetilde{\mathcal{A}} := \sum_{i=1}^r \widetilde{\mathcal{A}}_i, \quad \text{where } \widetilde{\mathcal{A}}_i = \text{mean}_{\mathcal{S}_\alpha}(\mathcal{A}_i^{(1)}, \dots, \mathcal{A}_i^{(M)}), \quad i = 1, \dots, r.$$

To guarantee that all points are comparable so that the Fréchet mean is well defined, we choose $\alpha = 1/\sqrt{k_1 + \cdots + k_d} - \sqrt{\epsilon}$ where ϵ is the machine precision of double-precision floating-point numbers. Proposition 4.15 then guarantees all points are comparable.

We reproduce the experiment from [CST23, Section 4.1], using the authors' code with the same parameter values and initialization seed. Figure 4.2 compares our estimate to their estimate and to the baseline of doing no consensus aggregation. The baseline, ZIPTF, consists of just doing one zero-inflated Poisson tensor factorization as described in [CST23, Section 3.2]. C-ZIPTF is the consensus aggregation of 20 ZIPTF's by Chafamo, Shanmugam, and Tokcan's proposed method of medians. F-ZIPTF is similarly the consensus aggregation of 20 ZIPTF's by our proposed method of Fréchet means. The Fréchet mean is approximated by successive geodesic interpolation, see for example [CV20; CV15; CV19; Che+16]. We conclude that the accuracy of both consensus variants is better than the baseline. There does not appear to be a meaningful difference between C-ZIPTF and F-ZIPTF. This nevertheless




Figure 4.2: We recreate the setup from [CST23, Section 4.1] and compare the Riemannian distance between the approximate decompositions and the ground truth.

validates the correctness and utility of connecting geodesics for the α -warped Segre manifold.

The Julia source code for the experiments is available publicly [Jac].

4.7 Conclusion and outlook

In summary, we have derived exponential and logarithmic maps for a one-parameter family of Riemannian metrics on the set of rank-one tensors. In this family, we identified several interesting members, including the Riemannian embedding in the Euclidean space of all tensors, as well as an interval of members that are geodesically connected.

The main tool that we used was the identification of the set of rank-one tensors locally as a *warped product*. This simplified the geometry enough that we were able to solve explicitly for the geodesics. Essentially, because of the warped product identification, we could reduce the geodesic equation to a one-dimensional ODE with an analytic solution. With the help of this identification, we were also able to derive closed-form expressions for the sectional curvatures. These expressions are very useful when approximating maps into rank-one tensors, because they determine the condition number for such an approximation, as we will discuss later in chapter 5. It is also this identification that was used to generalize to the one-parameter family.

The warped product identification works locally, but globally, the warped product is a normal Riemannian covering of rank-one tensors. We made this distinction more clear by introducing the *pre-Segre manifold*. Understanding this proved crucial in determining which points can be connected by geodesics, and which can not. We moreover provided a concrete algorithm, the *matchmaking algorithm*, for determining the shortest geodesics between compatible points.

We also mention that one can not easily identify tensors with higher ranks as warped products. Thus the results in this chapter can not easily be generalized to higher ranks.

Geodesics are ubiquitous in numerical analysis and numerical linear algebra, and in many situations additions and multiplications can be replaced by geodesic operations to generalize numeric methods from linear spaces to manifolds. We presented one such example application, consensus aggregation, where we used both the exponential and logarithmic map to approximate the Karcher mean of a set of rank-one terms. This is useful if you have several different decompositions of the same tensor, and want to combine them in a geometrically sound way.

There are many other natural applications for this work. For example Swijsen, Van der Veken, and Vannieuwenhoven [SvV22] use geodesics on the set of rank-one tensors to find low-CP-rank approximations. One could similarly use our one-parameter generalized geodesics for the same purpose, and get a one-parameter set of low-CP-rank approximation algorithms. Another natural application that we didn't talk about would be to solve manifold ODEs. It is straightforward to define a geodesic version of the Euler method [GOC25] to integrate an ODE intrinsically on the manifold of rank-one tensors.

Our work builds on a previous article by Swijsen, Vannieuwenhoven, and Van der Veken [SvV22]. Recall from the start of this chapter that this work was a collaboration between several people. There, it is also explained what precisely were my contributions.

4.8 Appendix: proofs of technical lemmas

4.8.1 Proof of lemma 4.1

Each ι_σ is an isometry because \mathcal{P}_α^k is a product manifold where ι_σ acts only on the spheres and the maps Id and $-\text{Id}$ are basic isometries of the sphere with the standard Euclidean inner product, and hence also with any scaled inner product.

Normality, i.e., $\otimes(p) = \otimes(q)$ implies there exists a σ such that $q = \iota_\sigma(p)$ [ONe83, Appendix A], follows from the multilinearity of the tensor product \otimes [Gre78, Chapter 1].

By the same arguments as [BBV19, Section 4.1], \otimes is a smooth covering map. Hence $\dim \mathcal{P}_\alpha^{\mathbf{k}} = \dim \mathcal{S}_\alpha^{\mathbf{k}}$ and the differential $d\otimes : T\mathcal{P}_\alpha^{\mathbf{k}} \rightarrow T\mathcal{S}_\alpha^{\mathbf{k}}$ is everywhere left-invertible.

The key part is showing that $\otimes : \mathcal{P}_\alpha^{\mathbf{k}} \rightarrow \mathcal{S}_\alpha^{\mathbf{k}}$ is Riemannian, i.e., the metric of $\mathcal{S}_\alpha^{\mathbf{k}}$ is the pushforward of the one of $\mathcal{P}_\alpha^{\mathbf{k}}$, which is a straightforward computation. Let $\dot{\mathbf{u}}_i \in T_{\mathbf{u}_i} S^{n_i-1} = \mathbf{u}_i^\perp$ and $\dot{\lambda} \in \mathbb{R}$ be arbitrary. By differentiating $\otimes : \mathcal{P}_\alpha^{\mathbf{k}} \rightarrow \mathcal{S}_\alpha^{\mathbf{k}}$, we get

$$\begin{aligned} & (d_{(\lambda, \mathbf{u}_1, \dots, \mathbf{u}_d)} \otimes)(\dot{\lambda}, \dot{\mathbf{u}}_1, \dots, \dot{\mathbf{u}}_d) \\ &= \dot{\lambda} \mathcal{U} + \lambda (\dot{\nu}_{k_1}(\dot{\mathbf{u}}_1) \otimes \mathbf{u}_2^{\otimes k_2} \otimes \dots \otimes \mathbf{u}_d^{\otimes k_d}) + \dots + \lambda (\mathbf{u}_1^{\otimes k_1} \otimes \dots \otimes \mathbf{u}_{d-1}^{\otimes k_{d-1}} \otimes \dot{\nu}_{k_d}(\dot{\mathbf{u}}_d)), \end{aligned}$$

where \mathcal{U} and $\nu_k(\mathbf{u})$ are as in (4.14) and (4.15). The α -warped metric on $\mathcal{S}_\alpha^{\mathbf{k}}$ satisfies

$$\varsigma_{\lambda \mathcal{U}}^\alpha (d \otimes (\dot{x}, \dot{\mathbf{u}}_1, \dots, \dot{\mathbf{u}}_d), d \otimes (\dot{y}, \dot{\mathbf{v}}_1, \dots, \dot{\mathbf{v}}_d)) = \dot{x} \dot{y} + (\alpha \lambda)^2 \langle \dot{\mathcal{U}}, \dot{\mathcal{V}} \rangle,$$

having dropped the subscript $(\lambda, \mathbf{u}_1, \dots, \mathbf{u}_d)$ of the differential and where

$$\dot{\mathcal{U}} = (\dot{\nu}_{k_1}(\dot{\mathbf{u}}_1) \otimes \mathbf{u}_2^{\otimes k_2} \otimes \dots \otimes \mathbf{u}_d^{\otimes k_d}) + \dots + (\mathbf{u}_1^{\otimes k_1} \otimes \dots \otimes \mathbf{u}_{d-1}^{\otimes k_{d-1}} \otimes \dot{\nu}_{k_d}(\dot{\mathbf{u}}_d)),$$

$$\dot{\mathcal{V}} = (\dot{\nu}_{k_1}(\dot{\mathbf{v}}_1) \otimes \mathbf{u}_2^{\otimes k_2} \otimes \dots \otimes \mathbf{u}_d^{\otimes k_d}) + \dots + (\mathbf{u}_1^{\otimes k_1} \otimes \dots \otimes \mathbf{u}_{d-1}^{\otimes k_{d-1}} \otimes \dot{\nu}_{k_d}(\dot{\mathbf{v}}_d)).$$

Recall from [Hac19] that the Euclidean inner product between rank-1 tensors satisfies $\langle \mathbf{u}_1 \otimes \dots \otimes \mathbf{u}_d, \mathbf{v}_1 \otimes \dots \otimes \mathbf{v}_d \rangle = \prod_{i=1}^d \langle \mathbf{u}_i, \mathbf{v}_i \rangle$. Hence, rank-1 tensors are orthogonal to one another if they are orthogonal in at least one factor. We compute that $\langle \nu_{k_i}(\dot{\mathbf{u}}_i), \mathbf{u}_i^{\otimes k_i} \rangle = 0$, while $\langle \nu_{k_i}(\dot{\mathbf{u}}_i), \nu_{k_i}(\dot{\mathbf{v}}_i) \rangle = k_i \langle \dot{\mathbf{u}}_i, \dot{\mathbf{v}}_i \rangle$ if $\dot{\mathbf{u}}_i, \dot{\mathbf{v}}_i \in T_{\mathbf{u}_i} S^{n_i}$. As a consequence, we can observe that the cross terms in $\langle \dot{\mathcal{U}}, \dot{\mathcal{V}} \rangle$ vanish and we are left with

$$\langle \dot{\mathcal{U}}, \dot{\mathcal{V}} \rangle = \sum_{i=1}^d \langle \dot{\nu}_{k_i}(\dot{\mathbf{u}}_i), \dot{\nu}_{k_i}(\dot{\mathbf{v}}_i) \rangle = \sum_{i=1}^d k_i \langle \dot{\mathbf{u}}_i, \dot{\mathbf{v}}_i \rangle.$$

This shows that $\varsigma_{\lambda \mathcal{U}}$ is the pushforward of $g_{\lambda, \mathbf{u}_1, \dots, \mathbf{u}_d}$ under \otimes , concluding the proof. \square

4.8.2 Proof of proposition 4.14

We want to maximize $f = \sum_{i=1}^d \sigma_i \Delta_i$ over $\sigma_i \in \{-1, 1\}$ with the constraint that $\sigma_1^{k_1} \dots \sigma_d^{k_d} = 1$.

We can assume without loss of generality that all k_i 's are odd. Indeed, even k_i do not impact the constraint, so we can choose the corresponding sign so that $|\Delta_i| = \sigma_i \Delta_i$, which is clearly optimal, also without constraint, because of the bilinear structure of f . The matching algorithm chooses these signs precisely in this way.

We can further assume without loss of generality that all Δ_i are nonzero, for otherwise we could take the unconstrained optimum (4.21) and simply swap the sign of a σ_j for which $\Delta_j = 0$ to obtain a feasible solution that is also globally optimal, which is what the matching algorithm does.

Let $(\varsigma_1, \dots, \varsigma_d)$ be a constrained optimizer of f that differs in the least number of indices from the unconstrained optimizer $(\sigma_1, \dots, \sigma_d)$ from (4.21). Denote the number of different indices by $p = \#\{i \in \{1, \dots, d\} \mid \sigma_i \neq \varsigma_i\}$. Then, there are three cases:

- $p = 0$. The unconstrained optimum given by (4.21) and chosen by the algorithm satisfies the constraint.
- $p = 1$. If f_0 denotes the value of f for the unconstrained optimum, such a constrained optimum has value $f_0 - 2|\Delta_i|$. Hence, i must correspond to the minimal $|\Delta_i|$, which the matching algorithm indeed selects by construction.
- $p \geq 2$. Let $i \neq j$ be any two indices for which $\varsigma_i \neq \sigma_i$ and $\varsigma_j \neq \sigma_j$. By changing the signs of ς_i and ς_j to match those of σ_i and σ_j , respectively, this increases the objective by $2|\Delta_i| + 2|\Delta_j|$. Since $|\Delta_i|, |\Delta_j| > 0$, this modified, feasible constrained point with $p - 2$ different indices has a strictly higher objective value than $(\varsigma_1, \dots, \varsigma_d)$, which was supposed to be the constrained minimizer with minimal p . This is a contradiction, so this case cannot occur.

This concludes the proof. □

4.8.3 Proof of (4.12)

Let $C = \frac{\lambda}{\alpha\lambda M}$, and recall that $B = \arctan(C)$ and $A = \lambda \cos(B) = \frac{\lambda}{\sqrt{1+C^2}}$. Define

$$z(t) := \frac{\lambda + \dot{\lambda}t}{\alpha\lambda Mt} = \frac{1}{\alpha Mt} + C, \quad \text{so that} \quad z'(t) = -\frac{1}{\alpha Mt^2}. \quad (4.22)$$

In this notation, (4.12) becomes $\alpha s(t) = \frac{\pi}{2} - \arctan(z(t))$, so that on the one hand

$$\alpha s'(t) = -\frac{z'(t)}{1+z^2(t)}.$$

Since $\sqrt{\dot{\lambda}^2 + (\alpha\lambda M)^2} = \alpha\lambda M\sqrt{1+C^2}$, (4.11) simplifies as follows:

$$\begin{aligned} \alpha s'(t) &= \cos^2(\alpha s(t) + B)A^{-1}\sqrt{\dot{\lambda}^2 + (\alpha\lambda M)^2} \\ &= \cos^2(\alpha s(t) + B)\alpha M(1+C^2) \\ &= \alpha M(1+C^2)\cos^2\left(\frac{\pi}{2} - (\arctan(z(t)) - \arctan(C))\right) \\ &= \alpha M(1+C^2)\sin^2(\arctan(z(t)) - \arctan(C)). \end{aligned}$$

Using basic trigonometric identities, it is easy to obtain the next identity:

$$\sin(\arctan(x) - \arctan(y)) = \frac{x-y}{\sqrt{1+x^2}\sqrt{1+y^2}}. \quad (4.23)$$

This allows us to conclude the proof because

$$\alpha s'(t) = \alpha M(1+C^2)\frac{(z(t)-C)^2}{(1+z^2(t))(1+C^2)} = \alpha M\frac{(\alpha M t)^{-2}}{1+z^2(t)} = -\frac{z'(t)}{1+z^2(t)}.$$

4.8.4 Proof of (4.13)

Using the same notation as in section 4.8.3, we get that $\alpha s(1) = \frac{\pi}{2} - \arctan(z(1))$, so that

$$\begin{aligned} \gamma(1) = \mu(s(1)) &= \frac{\lambda \cos(B)}{\cos(\alpha s(1) + B)} = \frac{\lambda}{\sqrt{1+C^2}} \cdot \frac{1}{\sin(\arctan(z(1)) - \arctan(C))} \\ &= \frac{\lambda}{\sqrt{1+C^2}} \frac{\sqrt{1+C^2}\sqrt{1+z^2(1)}}{z(1)-C} \\ &= \alpha\lambda M\sqrt{1+z^2(1)}, \end{aligned}$$

where we used (4.23) in the penultimate and (4.22) in the final equality. Since $z^2(1) = (\lambda + \dot{\lambda})^2(\alpha\lambda M)^{-2}$, the proof is concluded.

5. *Approximation theory on manifolds*

This chapter is based on the paper “Approximating maps into manifolds with lower curvature bounds” [Jac+25], where I am the main author. The paper is published in BIT Numerical Mathematics journal. I stated and proved the novel theorem and propositions, implemented the code and designed and ran the numerical experiments, drafted most of the text, incorporated feedback from the other authors, and drew the figures. Compared to the paper, we extend the discussion in section 5.2 of classical approximation theory for real-valued functions. We also discuss in section 5.4 more in-depth those manifolds most relevant to numerical analysis applications, along with the efficient evaluation and approximations of their exponential maps.

5.1 Introduction

Approximation theory is concerned with approximating functions with simpler functions. Classic approximation theory studies schemes for approximating real-valued functions. For example, univariate schemes like interpolation or regression [CC04; Tre19], but also multivariate schemes like hyperbolic crosses, sparse grids, greedy approximation, and moving least squares [Tem18; Wen04]. More recently, multivariate function approximation methods based on tensor decompositions have been extensively investigated [BM02; BM05; BEM16; DKS21; GSS23; HT17; HN18; SSK24; SK23]. All of these schemes are naturally extended to approximate vector-valued functions by applying them component-wise.

Many interesting functions arising in applications have some further structure. However, naively applying aforementioned classic approximation schemes fails to preserve such structures. To remedy this, we introduce a scheme to approximate maps whose target set is a *Riemannian manifold*. Our scheme works by pulling back the approximation problem to the tangent space at a

chosen point in the manifold. In this way, the problem is reduced again to the classic function approximation problem between vector spaces. Consequently, our scheme can leverage all of the previously mentioned classic techniques.

Some examples from applications of structured data that can be described as points on a Riemannian manifold are orthogonal matrices from motion tracking, matrix decompositions, differential equations, statistics, and optimization [AMS08; BGW15; CV19; Zim20; ZB24]; symmetric positive definite matrices and tensors from diffusion tensor imaging, materials modeling, computer vision, and statistics [Che+16; Moa06; Pen20; PFA06; TP21; TP23]; fixed-rank matrices and tensors from compression, tensor completion, and dynamic low-rank approximation [KL07; OT10; Ose11; SvV22; VAV12]; linear subspaces from Krylov subspace methods and subspace tracking [BGW15; SSK20; Yan95; ZH22]; and sequences of nested linear subspaces from principal component analysis [BP23; MCB24]. For instance, the present work was conceived to address difficulties encountered in the context of reduced order models of a microstructure materials engineering problem in [De +22].

We propose to approximate a map $f: R \rightarrow M$, where R is a set and M is a Riemannian manifold, by a natural, generally applicable three-step template:

1. Choose a point $p \in M$ from which to linearize the approximation problem.
2. Pull f back to the tangent space $T_p M$ using a *normal coordinate chart* [Lee18], and approximate the resulting map $g := \log_p \circ f: R \rightarrow T_p M$ by $\hat{g}: R \rightarrow T_p M$ using an arbitrary approximation scheme for maps into vector spaces.
3. Push \hat{g} forward with \exp_p to yield $\hat{f} := \exp_p \circ \hat{g}: R \rightarrow M$.

While the proposed template is a classic approach, the main challenge is establishing a bound on the approximation error, as any discrepancy between g and \hat{g} will be propagated through the manifold exponential. This requires understanding how the endpoint of a geodesic, which is a solution of a particular initial value problem on a manifold, changes as the initial conditions of this differential equation are varied. Our main contribution, theorem 5.5, consists of showing that, in this framework, an exact error analysis in principle only requires knowledge of a lower bound on the sectional curvature.

Alternatively, if \exp_p and \log_p are not known explicitly or are expensive to compute, a *retraction* [AMS08; Bou23] may be used instead. In corollary 5.5.1, we show how the error analysis can be extended to this case as well.

Algorithm 1 in section 5.5 is a detailed description of a concrete algorithm that results from applying this template to maps with domain $R = [-1, 1]^m$.

We implement it in a Julia package `ManiFactor.jl`. Figure 5.1 is a small example of approximating the map $f: [-1, 1]^2 \rightarrow S^2$ into the sphere, defined by stereographic projection of $(x, y) \mapsto (x^2 - y^2, 2xy)$.




Figure 5.1: Using different number of samples, we approximate $f: [-1, 1]^2 \rightarrow S^2$. Let Γ be a 10×10 grid on $[-1, 1]^2$. Solid lines show $f(\Gamma)$ and dashed lines show $\hat{f}(\Gamma)$.

It is natural to consider extending this scheme to multiple tangent spaces as well. Partly since it might guarantee a smaller error, but also since, if we are given a set of sample points, their image might not fit in a single normal coordinate chart. This extension was subsequently studied by Wang, Vandebril, Van der Veken, and Vannieuwenhoven [Wan+25].

5.1.1 Recent work

A lot of research is being done on approximating maps into manifolds. For example Dyn and Sharon [DS17a; DS17b] do manifold curve fitting by adapting subdivision schemes. Hardering and Wirth [HW21] and Heeren, Rumpf, and Wirth [HRW19] and Zhang and Noakes [ZN19] consider manifold versions of cubic splines. Bergmann and Gousenbourger [BG18], and Gousenbourger, Massart, and Absil [GMA19] define manifold Bézier curves. Sharon, Cohen, and Wendland [SCW23], Petersen and Müller [PM19], and Grohs, Sprecher, and Yu [GSY17] generalize moving least squares to manifolds. Cornea, Zhu, Kim, and Ibrahim [Cor+16], Hinkle, Fletcher, and Joshi [HFJ14], and Fletcher [Tho13] do regression of data on manifolds.

Also similar problems receive attention. The *Hermite interpolation problem* is to find a map that has given values and given derivatives in a set of sample points. Such a map can also be used for approximation. Séguin and Kressner [SCK24], Vardi, Dyn, and Sharon [BDS24], Zimmermann [Zim20; ZH22], Zimmermann and Bergmann [ZB24], and Moosmüller [Moo16; Moo17], solve Hermite problems for maps into manifolds.

There are also efforts to adapt statistics tools to manifolds. We only mention a few here. Gebhardt, Schubert, and Steinbach [GSS23], Curry, Marsland, and McLachlan [CMM19], Lazar and Lin [LL17], and Chakraborty, Seo, and Vemuri [CSV16] study how principal component analysis can be extended to different manifolds, a tool known as *principal geodesic analysis*. Pennec [Pen20] discusses these techniques specifically in the context of *diffusion tensor imaging*. Diepeveen, Chew, and Needell [DCN25] develop curvature corrected versions of tensor decompositions to compress manifold data. Diepeveen [Die24] extends several other common data analysis tools to manifolds.

Many of these works, for example [CSV16; Die24; DCN25; GMA19; SCK24; Zim20; ZH22; ZB24], fit into our template from the previous subsection, and our error analysis applies directly to them.

5.2 Classical approximation theory

We make no claim to overview the whole topic, but rather just aim to present some of the most useful and standard techniques to approximate univariate functions, along with some recent efforts to generalize these to the multivariate setting.

5.2.1 Univariate approximation theory

We consider the problem of expressing a real univariate function, h , approximately in terms of simpler functions. In practice, *simpler function* typically means something that can be evaluated using the four *basic operations* addition, subtraction, multiplication, and division—like a polynomial or rational function. The error that such an approximation introduces is well understood in terms of h 's smoothness, which is often encoded as some vector space norm. We will use the term *Banach space*, which formally is a complete normed vector space, and typically is some class of functions with some linear combination of L^p and max norms of different orders of derivatives.

One of the best ways to approximate functions with polynomials is with *Chebyshev interpolation*.

Proposition 5.1 (Trefethen [Tre19, Theorem 7.2]). *Let $h: [-1, 1] \rightarrow \mathbb{R}$ be $\nu \geq 1$ times differentiable such that $\|h^{(\nu)}\|_{\max} \leq V$ for some constant V . Let p_n be the degree n polynomial interpolating h in the Chebyshev nodes $t_k = \cos(k\pi/n)$,*

$k = 0, \dots, n$. Then, for $n \geq \nu$,

$$\|h - p_n\|_{\max} \leq \frac{4V}{\pi\nu(n - \nu)^\nu}. \tag{5.1}$$

We say that Chebyshev interpolating differentiable functions has *algebraic convergence*.

Proposition 5.2 (Trefethen [Tre19, Theorem 8.2]). *Let $h: [-1, 1] \rightarrow \mathbb{R}$ be analytic with an analytic continuation inside the Bernstein ellipse E_ρ , defined as the interior of the curve $\gamma(\theta) = (\rho e^{i\theta} + \rho^{-1} e^{-i\theta})/2$. Moreover, assume $\|h\|_{\max} \leq M$ on E_ρ for some constant M . Again, let p_n be the degree n polynomial interpolating h in the Chebyshev nodes. Then*

$$\|h - p_n\|_{\max} \leq \frac{4M\rho^{-n}}{\rho - 1}. \tag{5.2}$$

We say that Chebyshev interpolating analytic functions has *exponential convergence*.

Proposition 5.1 is tight in the sense that if h is ν times differentiable, but not $\nu + 1$ times differentiable, then the error will asymptotically decay as $n^{-\nu}$. Similarly for proposition 5.2, if the approximation error is majorized by ρ^{-n} , then f is analytic at least in E_ρ . Chebyshev interpolants are also tight in another sense. Let p_n^* be the degree n polynomial that minimizes the error $\|h - p_n^*\|_{\max}$. Then the Chebyshev interpolant is *near-best* [Tre19, Theorem 16.1] in the sense that

$$\|h - p_n\|_{\max} \leq \left(2 + \frac{2}{\pi} \log(n + 1)\right) \|h - p_n^*\|_{\max}. \tag{5.3}$$

Next, we note that p_n interpolates h in *prescribed points*. Moreover, the approximation scheme $\mathcal{I}_n: h \mapsto p_n$ is \mathbb{R} -linear. This is important for two reasons. First, it allows us to easily analyze error propagation. If there is some error ϵ in the evaluation of h , then this error is at worst multiplied by the operator norm Λ_n of \mathcal{I}_n . Λ_n is also called the *Lebesgue constant*, and has a modest numerical value of $\Lambda_n = \frac{2}{\pi} \log(n + 1) + 0.52125\dots$. Second, it allows us to build multivariate approximation schemes by taking tensor products, something we make precise in the next subsection.

We also mention that the polynomial coefficients of a Chebyshev interpolant can be computed efficiently from the evaluations $h(t_k)$. Using a discrete Fourier transform, it can be done in $\mathcal{O}(n \log n)$ basic operations [AF68, Method I]. Evaluating the polynomial is then $\mathcal{O}(n)$ basic operations. Alternatively, if

we only want to evaluate the approximation once, we can use the *barycentric formula*. It is an explicit expression for $p_n(x)$ in terms of h that only requires $\mathcal{O}(n)$ basic operations [BT04].

A robust way to approximate functions with rational functions is with *Padé approximants*. The type (m, n) Padé approximant of h can informally be defined as the quotient $r_{mn} = p_m/q_n$ satisfying

$$h(x) - r_{mn}(x) = \mathcal{O}(x)^{m+n+1}. \quad (5.4)$$

Such an approximation is useful typically when we have some expression for h so that we can obtain its Taylor coefficients. For example, we used Padé approximation in section 3.2.6 to approximate $(e^x - 1)/x$. The Padé approximant can be computed efficiently from the Taylor coefficients using an SVD of a certain Toeplitz matrix with the coefficients as entries [GGT13].

Unlike Chebyshev interpolation, the approximation scheme $\mathcal{P}_{mn}: h \mapsto r_{mn}$ is nonlinear. There is thus no analogue of the Lebesgue constant, and no canonical way to build multivariate schemes.

5.2.2 Multivariate approximation theory

Starting from some univariate approximation scheme, there are several approaches to generalize to the multivariate setting. One main approach is to *tensorize* a linear univariate scheme.

Recall the definition of tensor product from proposition 2.22. If $\mathcal{P}(\mathbb{R} \rightarrow \mathbb{R})$ denotes the space of real univariate polynomials, then pointwise multiplication

$$\begin{aligned} \phi: \mathcal{P}(\mathbb{R} \rightarrow \mathbb{R}) \times \cdots \times \mathcal{P}(\mathbb{R} \rightarrow \mathbb{R}) &\rightarrow \mathcal{P}(\mathbb{R}^m \rightarrow \mathbb{R}), \\ (p_1(x_1), \dots, p_m(x_m)) &\mapsto p_1(x_1) \cdots p_m(x_m) \end{aligned} \quad (5.5)$$

is a tensor product. The linear closure of the ϕ 's image is the whole of $\mathcal{P}(\mathbb{R}^m \rightarrow \mathbb{R})$ since any polynomial is a sum of monomials. We have thus argued that

$$\mathcal{P}(\mathbb{R} \rightarrow \mathbb{R})^{\otimes m} = \mathcal{P}(\mathbb{R}^m \rightarrow \mathbb{R}). \quad (5.6)$$

For the purposes of this subsection, and because we are now dealing with function spaces rather than \mathbb{R}^n , we slightly modify the meaning of the tensor product. We do not just want the tensor product of $A \otimes B$ of two Banach spaces to be a linear space, rather we want to again be a *Banach space*. We refer

to Ryan [Rya02, section 2] for details, and just note here that such a product exists.

We can infer a lot from the polynomial example. Since polynomials are dense in many function spaces, and since Banach spaces are complete, taking the closure of (5.6) directly yields for example

$$C(K \rightarrow \mathbb{R})^{\otimes m} = C(K^m \rightarrow \mathbb{R}), \tag{5.7}$$

$$C^\nu(K \rightarrow \mathbb{R})^{\otimes m} = C^\nu(K^m \rightarrow \mathbb{R}), \tag{5.8}$$

$$L^p(K \rightarrow \mathbb{R})^{\otimes m} = L^p(K^m \rightarrow \mathbb{R}). \tag{5.9}$$

Here, $K \subset \mathbb{R}$ is compact, C is equipped with the max norm, C^ν is equipped with the Sobolev norm $\sum_{i=1}^\nu \|h^{(i)}\|_{\max}$, and L^p is equipped with the L^p -norm.

Now, the idea is to view univariate approximation schemes as bounded linear maps between Banach spaces. In section 5.2.1, we saw that degree- n Chebyshev interpolation,

$$\begin{aligned} \mathcal{I}_n : C^\nu([-1, 1] \rightarrow \mathbb{R}) &\rightarrow C^\nu([-1, 1] \rightarrow \mathbb{R}), \\ h &\mapsto p_n, \end{aligned} \tag{5.10}$$

is linear and bounded by the Lebesgue constant Λ_n . Thus, by (5.8),

$$\mathcal{I}_{n_1} \otimes \cdots \otimes \mathcal{I}_{n_m} : C^\nu([-1, 1]^m \rightarrow \mathbb{R}) \rightarrow C^\nu([-1, 1]^m \rightarrow \mathbb{R}), \tag{5.11}$$

is a multivariate approximation scheme!

The error analysis of such schemes goes back to Schultz [Sch69].

Theorem 5.3 (Generalization of Schultz [Sch69, Theorem 2.1]). *Let H_1, \dots, H_m be Banach spaces of functions of type $\mathbb{R} \rightarrow \mathbb{R}$. Assume $g \in H_1 \otimes \cdots \otimes H_m$ is a multivariate function of type $\mathbb{R}^m \rightarrow \mathbb{R}$. That is, for all $k = 1, \dots, m$,*

$$h_k(t) := g(x_1, \dots, x_{k-1}, t, x_{k+1}, \dots, x_m) \in H_k \tag{5.12}$$

for all x_1, \dots, x_m . Moreover, for each k , let $\mathcal{F}_k : H_k \rightarrow H_k$ be a \mathbb{R} -linear approximation scheme. Assume \mathcal{F}_k has max error ϵ_k when applied to h_k , and let Ξ_k be its operator max norm. Finally, define $\hat{g} = (\mathcal{F}_1 \otimes \cdots \otimes \mathcal{F}_m)(g + e)$, where e represents some sort of measurement noise. Then

$$\|g - \hat{g}\| \leq \epsilon_1 + \Xi_1 \epsilon_2 + \cdots + \Xi_1 \cdots \Xi_{m-1} \epsilon_m + \Xi_1 \cdots \Xi_m \|e\|. \tag{5.13}$$

Proof.

$$\|\widehat{g} - g\| = \|(\mathcal{F}_1 \otimes \cdots \otimes \mathcal{F}_m)(g + e) - g\| \quad (5.14)$$

$$\leq \|(\mathcal{F}_1 \otimes \cdots \otimes \mathcal{F}_m)g - g\| + \|(\mathcal{F}_1 \otimes \cdots \otimes \mathcal{F}_m)e\| \quad (5.15)$$

$$\leq \|(\mathcal{F}_1 \otimes \cdots \otimes \mathcal{F}_m)g - g\| + \Xi_1 \cdots \Xi_m \|e\|. \quad (5.16)$$

The first term of (5.16) may be bounded by induction in m ,

$$\begin{aligned} & \|(\mathcal{F}_1 \otimes \cdots \otimes \mathcal{F}_m)g - g\| \\ &= \|(\mathcal{F}_1 \otimes \cdots \otimes \mathcal{F}_m - \mathcal{F}_1 \otimes \text{id} \otimes \cdots \otimes \text{id} + \mathcal{F}_1 \otimes \text{id} \otimes \cdots \otimes \text{id})g - g\| \end{aligned} \quad (5.17)$$

$$\leq \|(\mathcal{F}_1 \otimes \cdots \otimes \mathcal{F}_m - \mathcal{F}_1 \otimes \text{id} \otimes \cdots \otimes \text{id})g\| + \|(\mathcal{F}_1 \otimes \text{id} \otimes \cdots \otimes \text{id})g - g\| \quad (5.18)$$

$$\leq \|(\mathcal{F}_1 \otimes \cdots \otimes \mathcal{F}_m - \mathcal{F}_1 \otimes \text{id} \otimes \cdots \otimes \text{id})g\| + \epsilon_1 \quad (5.19)$$

$$= \|\mathcal{F}_1 \otimes (\mathcal{F}_2 \otimes \cdots \otimes \mathcal{F}_m - \text{id} \otimes \cdots \otimes \text{id})g\| + \epsilon_1 \quad (5.20)$$

$$\leq \Xi_1 \|(\text{id} \otimes \mathcal{F}_2 \otimes \cdots \otimes \mathcal{F}_m)g - g\| + \epsilon_1. \quad (5.21)$$

The base case, $m = 0$, has error 0. \square

Consider the tensorized Chebyshev interpolation that we introduced in (5.11). Proposition 5.2 shows that tensorized Chebyshev realizes the conditions in theorem 5.3. We have the following result.

Proposition 5.4. *Let $g: [-1, 1]^m \rightarrow \mathbb{R}$. Assume g is analytic in its k th argument with an analytic continuation inside the Bernstein ellipse with radius ρ_k , bounded by C_k . Define $\widehat{g} = (\mathcal{I}_{N_1} \otimes \cdots \otimes \mathcal{I}_{N_m})(g + e)$, where e represents some evaluation error. Then $\widehat{g}: [-1, 1]^m \rightarrow \mathbb{R}$ satisfies*

$$\begin{aligned} \|g - \widehat{g}\|_\infty &\leq \frac{4C_1}{(\rho_1 - 1)\rho_1^{N_1}} + \frac{4\Lambda_{N_1}C_2}{(\rho_2 - 1)\rho_2^{N_2}} + \cdots + \frac{4\Lambda_{N_1} \cdots \Lambda_{N_{m-1}}C_m}{(\rho_m - 1)\rho_m^{N_m}} \\ &\quad + \Lambda_{N_1} \cdots \Lambda_{N_m} \|e\|_\infty, \end{aligned} \quad (5.22)$$

where $\Lambda_N \leq \frac{2}{\pi} \log(N + 1) + 1$ and $\|e\|_\infty$ is the max norm of e . Especially, if $N_1 = \cdots = N_m = N$, $\rho_1 = \cdots = \rho = \rho$, and $C_1 = \cdots = C_m = C$, then

$$\|g - \widehat{g}\|_\infty \leq \frac{4(\Lambda_N^m - 1)C}{(\rho - 1)\rho^N(\Lambda_N - 1)} + \Lambda_N^m \|e\|_\infty. \quad (5.23)$$

In particular, tensorized Chebyshev approximation yields quasi-optimal approximations of continuous functions by bounded-degree polynomials [Mas80].

While these error bounds have been known since at least the 60s, tensorized approximants have not been practical until somewhat recently. The reason is essentially a curse of dimension. To construct \hat{g} in practice, we must use some decomposition of the evaluation tensor $G_{i_1 \dots i_m} = g(t_{i_1}, \dots, t_{i_m})$. Consider, for example, a *Tucker decomposition* $G = (U^{(1)} \otimes \dots \otimes U^{(m)}) \cdot C$, where $C \in \mathbb{R}^{r_1 \times \dots \times r_m}$ and each $U^{(k)} \in \mathbb{R}^{N_k \times r_k}$, $k = 1, \dots, m$, is a matrix. Given such a decomposition, we can construct \hat{g} as

$$\hat{g}(x) = (\mathcal{I}_{N_1} \otimes \dots \otimes \mathcal{I}_{N_m})((U^{(1)} \otimes \dots \otimes U^{(m)}) \cdot C)(x) \tag{5.24}$$

$$= \sum_{a_1=1}^{r_1} \dots \sum_{a_m=1}^{r_m} C_{a_1 \dots a_m} h_1^{a_1}(x_1) \dots h_m^{a_m}(x_m), \tag{5.25}$$

where $h_k^a \in \mathbb{R}^{N_k}$ is the univariate Chebyshev interpolant constructed from the a th column of $U^{(k)} \in \mathbb{R}^{N_k \times r_k}$. Each factor is thus approximated independently.

The problem of efficient multivariate approximation is in this way linked to the problem of efficient tensor decompositions. We discussed decomposition methods in sections 2.5.3 to 2.5.5. We mention a few recent approaches here.

Dolgov, Kressner, and Strössner [DKS21] implement tensorized Chebyshev for trivariate functions. Their scheme is precisely the special case of (5.25) when $m = 3$, with an HOSVD to decompose the evaluation tensor. Their theorem 2.1 and lemma 2.2 follow from setting $m = 3$ in proposition 5.4. At the end of their paper, they note that the extension of their approximation scheme to higher orders is theoretically straightforward, but practically infeasible due to the exponential increase in sampling points. Another similar approach, but that does not need the full evaluation tensor, is presented by Strössner, Sun, and Kressner [SSK24]. They implement tensorized Chebyshev for arbitrary order multivariate functions, with TT-cross to decompose the evaluation tensor, and are in this way able to interpolate benchmark functions in hundreds of variables!

There are also different approaches that, rather than interpolate an evaluation tensor, define continuous versions of CUR, TT-cross, etc. Townsend and Trefethen [TT13] use a continuous analogue of submatrix selection to construct a continuous pseudoskeleton for approximating bivariate functions. Hashemi and Trefethen [HT17] similarly apply continuous pseudoskeletons to decompose trivariate functions by fixing one variable at the time. Gorodetsky, Karaman, and Marzouk [GKM19] replace the discrete pseudoskeletons in TT-cross with continuous versions to yield a continuous tensor train decomposition.

We haven't talked about the error analysis of these second kind of multivariate approximation schemes. We instead refer to Griebel and Harbrecht's [GH23] thorough analysis of both continuous Tucker approximation and continuous tensor train approximation on tensor product *Sobolev spaces*.

5.3 Approximation theory on manifolds

The distance between two points on a Riemannian manifold depends on the geometry of that manifold. For example, the distance between two orthogonal matrices A and B in the space of orthogonal matrices will always be greater than or equal to the distance between A and B in the ambient space of matrices. When approximating maps into Riemannian manifolds it is natural to measure the approximation error *intrinsically* in the manifold, rather than in some *extrinsic* ambient space in which the manifold could be embedded. Moreover, for embedded manifolds with large codimension, such as low-rank matrices or tensors, it can be inefficient or even infeasible to measure the error in the ambient space. The next result thus bounds the intrinsic error of the approximation template from section 5.1.

Theorem 5.5. *Let M be a Riemannian manifold with sectional curvature bounded from below by some constant H and let $f: R \rightarrow M$, where R is a set whose image fits in a single geodesically convex normal coordinate chart $S \subset T_p M$ around $p \in M$ and assume that there exists some constant σ such that $\|X\|_p \leq \sigma$ for all $X \in S$. Let*

$$g = \log_p \circ f: R \rightarrow S \quad \text{and} \quad \hat{f} = \exp_p \circ \hat{g}: R \rightarrow M, \quad (5.26)$$

where $\hat{g}: R \rightarrow S$ is an approximation of g such that

$$\|g(x) - \hat{g}(x)\|_p \leq \epsilon \quad (5.27)$$

for all $x \in R$. Then, the distance between $f(x)$ and $\hat{f}(x)$ on M obeys

$$d_M(f(x), \hat{f}(x)) \leq \epsilon \quad \text{if } H \geq 0, \quad (5.28)$$

$$d_M(f(x), \hat{f}(x)) \leq \epsilon + \frac{2}{\sqrt{|H|}} \operatorname{arcsinh} \left(\frac{\epsilon \sinh(\sigma \sqrt{|H|})}{2\sigma} \right) \quad \text{if } H < 0, \quad (5.29)$$

for all $x \in R$.

Remark. We can in turn upper bound (5.29) by

$$d_M(f(x), \hat{f}(x)) \leq \epsilon + \frac{2}{\sqrt{|H|}} \log \left(\frac{\epsilon \exp(\sigma \sqrt{|H|})}{2\sigma} + 1 \right). \quad (5.30)$$

The assumptions in theorem 5.5 about geodesic convexity, bounded curvature, and bounded normal coordinate chart may seem very specific, but we note that they are always satisfied in some neighbourhood of p since curvature is continuous. To formulate explicit bounds, we however need explicit expressions for H . Therefore, we list several standard manifolds along with explicit lower bounds for their sectional curvature in section 5.4.

Proof of theorem 5.5. Let $x \in R$. The proof can be summarized as applying Toponogov's theorem to the geodesic triangle $(f(x), \widehat{f}(x), p)$, as visualized in figure 5.2.




Figure 5.2: Geodesic triangles defined by $p, f,$ and \widehat{f} . The map from M to N is $\exp_{N,q} \circ \log_{M,p}$.

Let N be the manifold of constant curvature H and let $q \in N$. Fix some isometry $T_p M \sim T_q N$ and let $\exp_{N,q}: S \rightarrow N$ be the manifold exponential on N . Define $h = \exp_{N,q} \circ g$ and $\widehat{h} = \exp_{N,q} \circ \widehat{g}$. By proposition 2.11,

$$d_M(f(x), \widehat{f}(x)) \leq d_N(h(x), \widehat{h}(x)). \tag{5.31}$$

Now our task is to bound the right-hand side of this inequality.

For brevity, from now on we drop the argument x of the functions $f, g,$ and h .

Case 1: $H \geq 0$. Since 0 lower bounds K we can, without loss of generality, let $H = 0$. Then $N = T_p M$ and so

$$d_N(h, \widehat{h}) = \|g - \widehat{g}\|_p \leq \epsilon. \tag{5.32}$$

In particular, the manifold exponential is a (weak) contraction.

Case 2: $H < 0$. Consider the geodesic triangle $(h(x), \widehat{h}(x), q)$. Define $c(t) = \cosh(t\sqrt{|H|})$ and $s(t) = \sinh(t\sqrt{|H|})$. Then lemma 2.13 says that

$$c(d(h, \widehat{h})) = c(\|g\|_p)c(\|\widehat{g}\|_p) - s(\|g\|_p)s(\|\widehat{g}\|_p)\cos(\angle(g, \widehat{g})). \tag{5.33}$$

Since \sinh is positive on \mathbb{R}^+ ,

$$s(\|g\|_p)s(\|\widehat{g}\|_p) \geq 0. \quad (5.34)$$

Hence, using the polarization identity,

$$\cos(\angle(g, \widehat{g})) = \frac{\langle g, \widehat{g} \rangle_p}{\|g\|_p \|\widehat{g}\|_p} = \frac{\|g\|_p^2 + \|\widehat{g}\|_p^2 - \|g - \widehat{g}\|_p^2}{2\|g\|_p \|\widehat{g}\|_p} \geq 1 - \frac{\|g - \widehat{g}\|_p^2}{2\|g\|_p \|\widehat{g}\|_p} \quad (5.35)$$

in (5.33) gives us

$$c(d(h, \widehat{h})) \leq c(\|g\|_p)c(\|\widehat{g}\|_p) - s(\|g\|_p)s(\|\widehat{g}\|_p) \left(1 - \frac{\|g - \widehat{g}\|_p^2}{2\|g\|_p \|\widehat{g}\|_p} \right) \quad (5.36)$$

$$= c(\|g\|_p - \|\widehat{g}\|_p) + \frac{\|g - \widehat{g}\|_p^2 s(\|g\|_p)s(\|\widehat{g}\|_p)}{2\|g\|_p \|\widehat{g}\|_p}. \quad (5.37)$$

Since \cosh is even and increasing on \mathbb{R}^+ , we may use the reverse triangle inequality to see that

$$c(d(h, \widehat{h})) \leq c(\|g - \widehat{g}\|_p) + \frac{\|g - \widehat{g}\|_p^2 s(\|g\|_p)s(\|\widehat{g}\|_p)}{2\|g\|_p \|\widehat{g}\|_p}. \quad (5.38)$$

Furthermore, noting that $\frac{\sinh x}{x}$ is increasing and positive on \mathbb{R}^+ , we have that

$$c(d(h, \widehat{h})) \leq c(\|g - \widehat{g}\|_p) + \frac{\|g - \widehat{g}\|_p^2 s(\sigma)^2}{2\sigma^2}. \quad (5.39)$$

This is equivalent to

$$1 + 2s\left(\frac{d(h, \widehat{h})}{2}\right)^2 \leq 1 + 2s\left(\frac{\|g - \widehat{g}\|_p}{2}\right)^2 + \frac{\|g - \widehat{g}\|_p^2 s(\sigma)^2}{2\sigma^2}, \quad (5.40)$$

so that

$$d(h, \widehat{h}) \leq 2s^{-1}\left(\sqrt{s\left(\frac{\|g - \widehat{g}\|_p}{2}\right)^2 + \frac{\|g - \widehat{g}\|_p^2 s(\sigma)^2}{4\sigma^2}}\right), \quad (5.41)$$

and since \sinh , $\operatorname{arcsinh}$, \cdot^2 , and $\sqrt{\cdot}$ are all increasing on \mathbb{R}^+ , we find

$$d(h, \widehat{h}) \leq 2s^{-1}\left(\sqrt{s\left(\frac{\epsilon}{2}\right)^2 + \frac{\epsilon^2 s(\sigma)^2}{4\sigma^2}}\right). \quad (5.42)$$

Lastly we will use that concave functions are *subadditive*, meaning $\sqrt{a^2 + b^2} \leq a + b$ and $\operatorname{arcsinh}(\sinh a + \sinh b) \leq a + b$ for nonnegative a and b . Thus

$$d(h, \widehat{h}) \leq \epsilon + 2s^{-1} \left(\frac{\epsilon s(\sigma)}{2\sigma} \right). \quad (5.43)$$

This concludes the proof. \square

5.3.1 Retractions

For many manifolds that appear in numerical analysis, we have closed form expressions for the exponential and logarithmic maps. They can also often be computed efficiently.

However, even if the exponential or logarithmic map are not known for a manifold, we might still be able to proceed by using a *retraction*. This is an injective map $r_p: S \rightarrow M$, for some open subset $S \subset T_p M$, that approximates the manifold exponential at p , and whose inverse approximates the logarithmic map. For a precise definition, see Absil, Mahony, and Sepulchre [AMS08, definition 4.1.1].

The approximant \widehat{g} of g is typically built from samples of g . If the samples are not exact however, then there is a priori no reason for the approximant to stay close to g . Thus, to study the error when using retractions, we need to introduce some additional assumption. We will therefore assume that the map $g \mapsto \widehat{g}$ is linear. This assumption is a balance between being specific enough for us to study and being generic enough to include common approximation schemes such as interpolation, projection, and least squares regression. It excludes adaptive approximation schemes, where the next sampling point can depend on the function value in the previous sampling point, such as for example AAA [NST18].

Corollary 5.5.1. *First, assume the same conditions as theorem 5.5. Then, assume that both the error between the retraction and the exponential map and the error between the inverse of the retraction and the logarithmic map are bounded:*

$$d_M(r_p(X), \exp_p(X)) \leq \zeta \quad \text{for all } X \in S, \quad (5.44)$$

$$\|\log_p(q) - r_p^{-1}(q)\|_p \leq \eta \quad \text{for all } q \in \exp_p(S). \quad (5.45)$$

Moreover, assume that the vector-valued approximation scheme $\mathcal{F}: g \mapsto \widehat{g}$ is \mathbb{R} -linear with operator max norm Λ . Now, instead of using $\widehat{f} = (\exp_p \circ \mathcal{F} \circ \log_p)(f)$,

define $\widehat{f} = (r_p \circ \mathcal{F} \circ r_p^{-1})(f)$. Then

$$d_M(f(x), \widehat{f}(x)) \leq \epsilon + \Lambda\eta + \zeta \quad (5.46)$$

if $H \geq 0$

$$d_M(f(x), \widehat{f}(x)) \leq \epsilon + \Lambda\eta + \zeta + \frac{2}{\sqrt{|H|}} \operatorname{arcsinh} \left(\frac{(\epsilon + \Lambda\eta) \sinh(\sigma\sqrt{|H|})}{2\sigma} \right) \quad (5.47)$$

if $H < 0$.

Proof. We can reduce to theorem 5.5 by letting $\widehat{g}_1 = (\mathcal{F} \circ r_p^{-1})(f)$, $\widehat{g}_2 = (\mathcal{F} \circ \log_p)(f)$, and noticing that

$$\|g - \widehat{g}_1\| \leq \|g - \widehat{g}_2\| + \|\widehat{g}_2 - \widehat{g}_1\| \leq \|g - \mathcal{F}(g)\| + \Lambda \|r_p^{-1}(f) - \log_p(f)\| \leq \epsilon + \Lambda\eta. \quad (5.48)$$

This concludes the proof. \square

If M is an embedded submanifold, one can construct a retraction around a point $p \in M$ by embedding the tangent space and then projecting from there back to M . Alternatively, if $\mu: T_p M \rightarrow M$ is a smooth chart around p , then $\mu \circ (d_p \mu)^{-1}$ is a retraction around p . If M is a quotient manifold, any retraction on the total space induces a retraction on M , and defining a retraction on the total space often amounts to approximating the matrix exponential. Absil et al [AMS08, Section 4.1] is an exposition on these different ways of retracting. See also the references in [AMS08, Section 4.10] on how to implement them efficiently for different manifolds. Absil and Oseledet [AO15] survey retractions on the manifold of fixed-rank matrices, for which the exponential or logarithmic maps are not known. The bottom row of their tables 1 and 2 list what we in corollary 5.5.1 call ζ .

5.3.2 Condition numbers

When $H \rightarrow 0$ from below, (5.29) tends to

$$d_M(f, \widehat{f}) \leq 2\epsilon, \quad (5.49)$$

which shows that it is not a tight bound. To explain why, first note that (5.35) is an equality when $\|g(x)\|_p = \|\widehat{g}(x)\|_p$, but we also used the reverse triangle inequality in (5.38) which is an equality when g and \widehat{g} are collinear. Furthermore,

in (5.43) we used subadditivity for the square root, which is an equality only when one of the terms is 0. However, for small curvatures and small ϵ , the two terms in (5.43) are approximately equal.

Even though the bound (5.29) is not tight, there is, in the following sense, no better bound.

Proposition 5.6. *Recall the notation of theorem 5.5. Assume that M has constant negative sectional curvature H and let \widehat{g} be an approximation to g such that the maximum error ϵ is attained at a point $x \in R$ with $\|\widehat{g}(x)\|_p = \|g(x)\|_p = \sigma$. Then*

$$d_M(f(x), \widehat{f}(x)) = \frac{2}{\sqrt{|H|}} \operatorname{arcsinh} \left(\frac{\epsilon \sinh(\sigma \sqrt{|H|})}{2\sigma} \right). \tag{5.50}$$

Proof. Consider again the proof of theorem 5.5. If $\|g\|_p = \|\widehat{g}\|_p$, then (5.35) is an equality, and the reverse triangle inequality is not required in (5.38), hence (5.40) reduces to

$$1 + 2s \left(\frac{d_M(f, \widehat{f})}{2} \right)^2 = 1 + \frac{\epsilon^2 s(\sigma)^2}{2\sigma^2}. \tag{5.51}$$

This is equivalent to what we wanted to prove. □

As an immediate consequence of theorem 5.5 we can obtain bounds on the propagation of a small perturbation $g - \widehat{g}$ through the manifold exponential \exp_p . The size of this output perturbation relative to the input perturbation is quantified by the *condition number*. Recall the definition of Rice’s condition number [Ric66] of a map $\phi: A \rightarrow B$ between metric spaces (A, d_A) and (B, d_B) :

$$\kappa[\phi](x) := \limsup_{\substack{x' \in A, \\ d_A(x, x') \rightarrow 0}} \frac{d_B(\phi(x), \phi(x'))}{d_A(x, x')}. \tag{5.52}$$

Recall R and S from theorem 5.5. Let $C(R, S)$ denote the vector space of continuous functions from R to S and define the *pushforward* $\phi: C(R, S) \rightarrow C(R, M)$ by $\phi(g) := \exp_p \circ g$. Further, define the metrics

$$d_{C(R, S)}(g, \widehat{g}) := \sup_{x \in R} \|g(x) - \widehat{g}(x)\|_p, \quad \text{and} \quad d_{C(R, M)}(f, \widehat{f}) := \sup_{x \in R} d_M(f(x), \widehat{f}(x)). \tag{5.53}$$

The condition number implies an asymptotically sharp error bound

$$d_{C(R,M)}(\exp_p \circ g, \exp_p \circ \hat{g}) \leq \kappa[\exp_p](g)d_{C(R,S)}(g, \hat{g}) + o(d_{C(R,S)}(g, \hat{g})) \quad (5.54)$$

$$= \kappa[\exp_p](g)\epsilon + o(\epsilon), \quad (5.55)$$

Proposition 5.7. *Recall f and g from the notation of theorem 5.5. The condition number κ of the map $\phi: g \mapsto f$ satisfies*

$$1 \leq \kappa \leq 1 + \frac{\sinh(\sigma\sqrt{|H|})}{\sigma\sqrt{|H|}} \quad \text{if } H < 0, \quad (5.56)$$

$$\kappa = 1 \quad \text{if } H \geq 0. \quad (5.57)$$

Furthermore, if M is a model manifold with constant negative sectional curvature H , then κ further satisfies

$$\frac{\sinh(\sigma\sqrt{|H|})}{\sigma\sqrt{|H|}} \leq \kappa. \quad (5.58)$$

Proof. The upper bound is a corollary of theorem 5.5. Moreover, $1 \leq \kappa$ since \exp_p is a radial isometry. When M is a model manifold, κ is lower bounded according to proposition 5.6. This concludes the proof. \square

Two main takeaways in this subsection are that the condition number grows exponentially in both σ and $\sqrt{|H|}$, but that choosing σ so that $\sigma\sqrt{|H|} \leq 1$ guarantees a reasonable condition number $\kappa \leq 1 + \sinh(1) \approx 2.17$. To control σ , one may have to partition the target manifold into several charts, so that each chart is only so big that a good condition number can be guaranteed. This may also involve using more samples, so that there are enough samples in each chart to construct a good approximation in the tangent space. See also Wang et al. [Wan+25] who define a *multiple tangent space model* for approximating maps into manifolds.

Zimmermann [Zim20, Theorem 3.1] derives a local version of proposition 5.7 using the Jacobi equation. Diepeveen et al. [DCN25, Theorem 3.4] arrive at a similar expression when they estimate tensor compression errors for data on locally symmetric manifolds, also using the Jacobi equation. The Jacobi equation however only encodes first-order information, and so an advantage of our approach is that we can derive *exact* error bounds with Toponogov's theorem.

5.3.3 Convexity radius and injectivity radius

Recall the definition of injectivity radius and Klingenberg's theorem from section 2.2.4.

Knowledge of the injectivity radius, along with upper curvature bounds, is very relevant, because it gives us information about how to choose the chart radius σ in theorem 5.5. For if i is the injectivity radius on M and C is an upper bounded for the sectional curvature, then $\sigma = \frac{1}{2} \min \{i, \pi/\sqrt{C}\}$ guarantees that $S = B_\sigma(p)$ is geodesically convex.

Cheeger and Ebin [CE08, theorem 5.14] prove this statement, originally due to Whitehead, using a *second variation* formula. If the injectivity radius is derived from Klingenberg's theorem using an upper curvature bound C , and if we are in the case $\pi/\sqrt{C} \leq l/2$, there is however a more visual interpretation of the statement. We can show that geodesic triangles with corners in S have a maximal "thickness", so that the whole triangle is contained in S . Essentially, the model manifold N is a sphere of radius $1/\sqrt{C}$, and σ is less than a quarter of the circumference, implying that the worst case scenario triangle still has its comparison triangle contained in one hemisphere T of the model manifold N , and hemispheres are geodesically convex. Now we can use the companion theorem, proposition 2.12, of Toponogov's that we discussed in section 2.2.3, to say that S is also geodesically convex. The argument is illustrated in figure 5.3.




Figure 5.3: The fattest possible comparison triangle is still contained in one hemisphere $T = B_\sigma(q)$.

We summarize what we have said about S 's geodesic convexity in a form that will be the most convenient for us to apply.

Proposition 5.8. *Let M be a compact Riemannian manifold with injectivity radius i and with sectional curvature bounded from above by some constant C ,*

and let $p \in M$. Let

$$\sigma = \min \left\{ \frac{\pi}{2\sqrt{C}}, \frac{i}{2} \right\}. \quad (5.59)$$

Then $B_\sigma(p)$ is geodesically convex.

Definition 5.9. The *convexity radius* is the maximal σ such that $B_\sigma(p)$ is geodesically convex.

Note that the upper curvature bound is important. Just setting $\sigma = i/2$ does not work. For example, Dibble [Dib17] shows that the ratio between the convexity radius and the injectivity radius can be made arbitrarily small.

5.4 Manifolds with bounded sectional curvature

In this section, we list Riemannian manifolds with explicit lower bounds for their sectional curvature K and whose manifold exponential and logarithm are known. For these manifolds, theorem 5.5 gives an explicit bound for the error on the manifold in terms of the error on the tangent space. While sectional curvature is always bounded locally, for many manifolds that are relevant for applications we can find explicit expressions for those local bounds. In many cases, the sectional curvature is even bounded globally, so that we can choose the same H for every chart.

When upper curvature bounds and injectivity and convexity radii are known, we mention those as well.

5.4.1 Lie groups

Given an inner product on its Lie algebra, one can extend it to a Riemannian metric on the Lie group by demanding invariance under left (or right) translation. Because of this invariance, the sectional curvature is the same at every point, and a bound for the sectional curvature at any one point is a global bound. Moreover, there is an explicit expression for the sectional curvature of a Lie group in terms of its Lie bracket.

Proposition 5.10 ([CE08, Proposition 3.18 (3)]). *Let $\langle \cdot, \cdot \rangle$ be a left-invariant inner product on the Lie group G . Let X, Y be left-invariant vector fields on G .*

Then

$$\begin{aligned}
 K(X, Y) = & \left(\|\text{ad}_X^* Y + \text{ad}_Y^* X\|^2 - \langle \text{ad}_X^* X, \text{ad}_Y^* Y \rangle - \frac{3}{4} \|[X, Y]\|^2 - \right. \\
 & \left. - \frac{1}{2} \langle [[X, Y], Y], X \rangle - \frac{1}{2} \langle [[Y, X], X], Y \rangle \right) / Q(X, Y), \quad (5.60)
 \end{aligned}$$

where $Q(X, Y) = \|X\|^2 \|Y\|^2 - \langle X, Y \rangle^2$ is the area of the parallelogram spanned by X and Y , and $\text{ad}_X = [X, \cdot]$ is the action of the adjoint representation.

We are interested in a specific class of Lie groups: matrix groups whose Lie algebra has the Euclidean inner product. There, we can use the inequality $\|[u, v]\| \leq \sqrt{2} \|u\| \|v\|$, which also implies $\|\text{ad}_u\|_{\text{op}} \leq \sqrt{2} \|u\|$, to show that¹

$$-\frac{11}{2} \leq K \leq 8. \quad (5.61)$$

From now on, this is the assumed metric on all Lie groups.

The most general example is the group, $\text{GL}(n)$, of invertible $n \times n$ matrices. Its metric is $\langle u, v \rangle_A = \text{tr} \left((A^{-1}u)^T A^{-1}v \right)$, where $u, v \in T_A \text{GL}(n) = \mathbb{R}^{n \times n}$. By Ado’s theorem [Lee13, theorem 8.49], all Lie groups are locally Lie subgroups of some $\text{GL}(n)$. Andruchow, Larotonda, Recht, and Varela [And+14] derive its geodesics. Recall the expression from (3.36), where we stated them for the right-invariant metric.

The compact Lie groups and their products with vector spaces are precisely the Lie groups that admit a *bi-invariant* metric [CE08, proposition 3.16 (6) and Proposition 3.34], i.e., a metric that is both left- and right-invariant. For such manifolds, there is a particularly simple expression for the sectional curvature.

Proposition 5.11 ([CE08, Corollary 3.19 (3)]). *Let $\langle \cdot, \cdot \rangle$ be a bi-invariant inner product on a Lie group G . Let X, Y be bi-invariant vector fields on G . Then*

$$K(X, Y) = \|[X, Y]\|^2 / Q(X, Y). \quad (5.62)$$

Especially

$$0 \leq K \leq 2. \quad (5.63)$$

The classification of compact Lie groups is a classic result. The real compact Lie groups are [Arv03, Theorem 2.17]

¹Without loss of generality, X and Y can be chosen to be orthonormal, so that $Q = 1$.

- tori $S^1 \times \cdots \times S^1$,
- orthogonal groups $O(n)$ and special orthogonal groups $SO(n)$,
- unitary groups $U(n)$ and special unitary groups $SU(n)$,
- spin groups $\text{Spin}(n)$,
- symplectic groups $\text{Sp}(n)$,
- compact exceptional Lie groups G_2 , F_4 , E_6 , and E_8 ,
- products, disjoint unions, and finite covers of the above.

Having a bi-invariant metric is also equivalent to the manifold exponential agreeing with the Lie group exponential [CE08, corollary 3.19]. Evaluating the manifold exponential on a compact matrix group is thus straightforward, it is the matrix exponential, $\text{mexp}(\cdot)$. The state of the art for computing the matrix exponential is the *scaling and squaring* method combined with a diagonal *Padé approximant* [Hig08, chapter 10]. The inverse, the matrix logarithm $\text{mlog}(\cdot)$, can similarly be computed via scaling and squaring, although some care has to be taken to choose the correct branch [Hig08, chapter 10]. The standard resource on this topic is Higham's book [Hig08]. We also mention Moler and Van Loan's classic survey [MV03].

The most important compact matrix group is probably $O(n)$. Its Lie algebra comprises skew-symmetric matrices, and so we can ask whether there is some way to take advantage of this structure to exponentiate more efficiently. For $O(3)$ specifically, there is a well-known algorithm [GX03], *Rodrigues' formula*,

$$\text{mexp}(A) = 1 + \frac{\sin \alpha}{\alpha} A + \frac{1 - \cos \alpha}{\alpha^2} A^2, \quad (5.64)$$

where $\alpha = \|A\|_{\mathbb{F}}/\sqrt{2}$, with natural applications to rigid body dynamics. Gallier and Xu [GX03] generalize the algorithm to $n \geq 4$ by block-diagonalizing A and applying Rodrigues' to each block. There is also the work by Cardoso and Leite [CL10] to improve the standard algorithm by optimizing the Padé approximant for skew-symmetric inputs. We also mention Deng, Absil, Gallivan, and Huang [Den+25] who consider efficient evaluations of the derivative $d(\text{mexp})$ for skew-symmetric input.

The injectivity radius for $O(n)$ is π , best explained in a blog post by Axen [Axe23]. His argument is straightforward enough that we summarize it here: a skew-symmetric matrix A is diagonalizable with purely imaginary diagonal, $A = PDP^{-1}$, $D = \text{diag}(i\theta_1, \dots, i\theta_n)$, and so the exponential of A is

invertible as long as all the $\theta_i \mapsto \exp(i\theta_i)$ are invertible. By proposition 5.8, the convexity radius for $O(n)$ is then at least $\pi/(2\sqrt{2})$.

The argument is easily extended to $GL(n)$. Recall again the $GL(n)$ geodesics from (3.36), and let $A \in \mathfrak{gl}(n)$ be a tangent at $1 \in GL(n)$. If A closes a geodesic loop, then

$$\text{mexp}(A^\top) \text{mexp}(A - A^\top) = 1 \tag{5.65}$$

$$\text{mexp}(A^\top) = \text{mexp}(A^\top - A) \tag{5.66}$$

{Use for example [Hig08, theorem 1.27]}

$$A = A - A^\top + \text{something skew-symmetric}, \tag{5.67}$$

and so A is skew-symmetric. Hence $A \in \mathfrak{o}(n)$. Since $O(n)$ is a totally geodesic submanifold of $GL(n)$, there are indeed geodesic loops with length 2π , but not shorter. So, again by proposition 5.8, the convexity radius for $GL(n)$ is at least $\pi/(4\sqrt{2})$.

While the matrix exponential is not a Riemannian exponential for more general matrix groups, it is still useful for the following reason: it is a retraction! The most straightforward way to see this is to use a definition of the Lie group exponential of tX as the one-parameter Lie subgroups in G that is tangent to X , and then to use that the matrix exponential is a Lie group exponential. For $GL(n)$, it can be seen directly using the BCH formula,

$$\text{mexp}(tX - tX^\top) \text{mexp}(tX^\top) = 1 + tX - tX^\top + tX^\top + \mathcal{O}(t^2) \tag{5.68}$$

$$= 1 + tX + \mathcal{O}(t^2). \tag{5.69}$$

For $O(n)$, there is also two classic retractions: the *QR* and *polar retractions*, based on the polar and QR decompositions respectively. If $A \in O(n)$ and $X \in T_A O(n)$, then the QR retraction is

$$R_A(X) = \text{qr}_1(A + X), \tag{5.70}$$

where qr_1 is the Q factor in a QR decomposition. The polar retraction is defined similarly.

5.4.2 Lie group quotients

Let $H \subset G$ be a Lie subgroup. By *Lie group quotient*, we mean the set of equivalence classes gH , equipped with a smooth structure by demanding that

$\pi: G \rightarrow G/H$ is a smooth submersion. This is well-defined by the smooth submersion characteristic property, proposition 2.5. If G has some Riemannian structure, we also demand that π is a Riemannian submersion. This is well-defined by proposition 2.19.

A *Riemannian homogeneous space* is a Riemannian manifold, along with some transitive action G that the metric is invariant under. Similarly to Lie groups, the curvature is then the same at every point. As discussed in sections 2.3.1 and 2.3.2, any Riemannian homogeneous space is a Lie group quotient. The converse is however not true, and Lie group quotients can only be Riemannian homogeneous if they are also reductive.

There is a formula for the sectional curvature of G/H in terms of the sectional curvature of G .

Proposition 5.12 ([ONe66]). *Let $\pi: M \rightarrow N$ be a Riemannian submersion. Let $X, Y \in T_p M$, and let $X', Y' \in T_{\pi(p)} N$ be horizontal vectors such that $d\pi X = X'$ and $d\pi Y = Y'$. Then*

$$K_N(X', Y') = K_M(X, Y) + \frac{3}{4} \|[X, Y]^V\|^2 / Q(X, Y), \quad (5.71)$$

where $[\cdot, \cdot]^V$ is the vertical component of $[\cdot, \cdot]$. Especially,

$$K_M \leq K_N, \quad (5.72)$$

and for Lie group quotients,

$$K_G \leq K_{G/H}. \quad (5.73)$$

If $\pi(p) = q \in N$, then the geodesics in N emanating from q are in one-to-one correspondence with horizontal geodesics in M emanating from p . In principle, this makes computing the exponential map for homogeneous spaces straightforward. However, if $\dim M \gg \dim N$, or N otherwise has some special structure, then we may be able to compute N 's exponential map much more efficiently.

The *Stiefel manifold*² $\text{St}(n, k) = \text{O}(n) / \text{O}(n - k)$ is the space of orthogonal $n \times k$ matrices. The horizontal space is just skew-symmetric matrices with zeros in the lower right $(n - k) \times (n - k)$ block. In Edelman, Arias, and Smith's [EAS98] foundational paper, they lay the groundwork for Riemannian

²The Riemannian homogeneous metric on the Stiefel is often called the *canonical metric* to differentiate it from its metric as an embedded submanifold of Euclidean space, which we discuss later.

optimization on the Stiefel manifold by showing that the matrix exponential of such a skew-symmetric matrix can be evaluated in $\mathcal{O}(nk^2)$ operations, compared to $\mathcal{O}(n^3)$ for the naive matrix exponential.

We now explain the fundamental difficulty of computing the logarithmic map on a homogeneous manifold $M = G/H$. Let $\mathfrak{g} = \mathfrak{h} \oplus \mathfrak{m}$ be the decomposition of G 's Lie algebra into vertical and horizontal components, and let $p = 1H$. Say we can compute $q = \exp_1(X)H$ efficiently when X is horizontal. In the case of the Stiefel manifold, Edelman et al. gives us an efficient algorithm to compute the first k columns of $\text{mexp}(X)$ when X is skew-symmetric with zeros in the lower left $(n - k) \times (n - k)$ block. However, given the equivalence class q , it is nontrivial to find the precise group element g where the horizontal geodesics intersects the fiber. For arbitrary $g' \in gH$, $X' = \log_1 g'$ is typically not horizontal, and so $\log_p q$ is *not* just $X' + \mathfrak{h}$. The problem is illustrated in figure 5.4.




Figure 5.4: γ is a horizontal geodesic in G emanating from 1. It cuts the fiber gH at a right angle in g .

One natural idea to compute the logarithmic map is to perform some sort of fixed point iteration on $\exp_1 \circ \text{proj}_{\mathfrak{m}} \circ \log_1$. Although, implemented naively, it requires doing a lot of expensive computations in the bigger space G . Zimmermann [Zim17] computes the Stiefel logarithm by fixed point iterating a smaller associated $2k \times 2k$ matrix and shows local exponential convergence.

A lot of attention has been given to the injectivity radius for the Stiefel manifold, especially in the context of the one-parameter family of metrics, introduced by Hüper, Markina, and Leite [KIS21], that includes both the canonical and Euclidean metrics. For the canonical metric, Zimmermann and Stoye [ZS25] show that $2\pi/\sqrt{5} \leq i$ using Klingenberg’s lemma. Their argument combines a bound for the sectional curvature [ZS25, theorem 8], $0 \leq K \leq 5/4$, and a lower bound for the length of geodesic loops [AM25, theorem 6.1],

$2\pi \leq l$, obtained from a similar diagonalization argument that we discussed in section 5.4.1. From this follows also that the convexity radius for the Stiefel is at least $\pi/\sqrt{5}$

The *Grassmann manifold* $\text{Gr}(n, k) = \text{O}(n) / (\text{O}(n - k) \times \text{O}(k))$ is the space of k -dimensional linear subspaces of \mathbb{R}^n . Bendokat, Zimmermann, and Absil [BZA24] consider efficient Riemannian exponentials and logarithms in the Grassmann manifold. The Grassmann manifold has its curvature curvature bounded by [BZA24; WC88]

$$0 \leq K \leq 2. \quad (5.74)$$

Its injectivity radius is $\pi/2$. From this follows that the convexity radius is at least $\pi/4$.

For the Stiefel and Grassmann manifolds, their metrics as Riemannian homogeneous spaces are called *canonical metrics* to distinguish them from their metrics as embedded submanifolds of $\mathbb{R}^{n \times n}$. With the canonical metric, they both have nonnegative curvature.

Let $0 = d_0 < d_1 < \dots < d_k = n$ be an increasing integer sequence. The *flag manifold* $\mathcal{F}_{(d_0, \dots, d_k)} = \text{O}(n) / (\text{O}(d_1 - d_0) \times \dots \times \text{O}(d_k - d_{k-1}))$ is the space of increasing sequences of linear subspaces in \mathbb{R}^n . It has nonnegative curvature, and its manifolds exponential can be computed efficiently by lifting geodesics to the Stiefel manifold.

$\sigma_B: A \mapsto BAB^\top$ defines a transitive action of $\text{GL}(n)$ on the space \mathcal{P}_n of *symmetric positive definite $n \times n$ matrices*. The *affine invariant metric* on \mathcal{P}_n is thus derived from the quotient $\mathcal{P}_n = \text{GL}(n) / \text{O}(n)$. See Bhatia [Bha07, Theorem 6.1.6] for an expression for its geodesics. It has nonpositive curvature.

Similarly to \mathcal{P}_n , σ_B defines a transitive action of $\text{GL}^+(n)$, the group of matrices with positive determinant, on the space $\text{SPSD}(n, k)$ of *rank- k positive semidefinite $n \times n$ matrices*. We have that

$$\text{SPSD}(n, k) = \text{GL}^+(n) / \begin{bmatrix} \text{SO}(k) & \mathbb{R}^{k \times (n-k)} \\ & \text{GL}^+(n-k) \end{bmatrix}. \quad (5.75)$$

Vandereycken, Absil, and Vandewalle [VAV12] use a right-invariant metric on $\text{GL}^+(n)$ to derive geodesics on $\text{SPSD}(n, k)$.

We also mention another metric that can be put on the space of rank- k positive semidefinite matrices. Let $\mathbb{R}_*^{n \times k}$ be the space of full rank $n \times k$ matrices. Any rank- k positive semidefinite matrix can be written as YY^\top , $Y \in \mathbb{R}_*^{n \times k}$. Such a representation is unique up to right multiplication by an

orthogonal matrix. Hence, we define $\mathcal{S}_+(n, k) = \mathbb{R}_*^{n \times k} / \mathrm{O}(k)$, where $\mathbb{R}_*^{n \times k}$ is equipped with the Euclidean metric. Although this manifold is not a Lie group quotient, a metric may still be induced by demanding that the quotient map is a Riemannian submersion. Massart and Absil [MA20] summarize its geometry. It has nonnegative curvature.

The manifolds we mentioned in sections 3.2 to 3.4 all fit into this framework as well, their sectional curvature lower bounded by $-11/2$ from (5.61).

We already mentioned that the exponential map on G induces an exponential map on G/H through π being a Riemannian submersion. Furthermore, a retraction on G induces a retraction on G/H through π being a smooth submersion. Since we also know at least one retraction on G , namely its Lie group exponential, we now know at least one retraction on any Lie group quotient. We discussed in section 3.2.6 how to efficiently compute the matrix exponential for quotients like (5.75) by using the low-rank structure of the tangents.

The QR and polar retractions on $\mathrm{O}(n)$ similarly induce retractions on $\mathrm{St}(n, k)$. Advantageously, the corresponding decompositions can be done efficiently, i.e. without having to do the decomposition in $\mathrm{O}(n)$. Computing the Q factor in the QR decomposition of a $n \times k$ matrix can for example be done in $2(n - k/3)k^2$ basic operations.

5.4.3 Embedded submanifolds in Euclidean space

We already mentioned that $\mathrm{O}(n)$ and $\mathrm{St}(n, k)$ have nonnegative curvature when equipped with the canonical metric. However, we can also equip them with the *Euclidean metric*, induced by embedding in $\mathbb{R}^{n \times n}$ and $\mathbb{R}^{n \times k}$ respectively. Zimmermann and Stoye [ZS25, theorem 10] show that $-1/2 \leq K \leq 1$ when using the Euclidean metric. Zimmermann and Hüper [ZH22] derive its exponential and logarithmic maps. We also mention that Mataire, Zimmermann, and Miolane [MZM25] extend the fixed point iteration method by Zimmermann [Zim17] that we mentioned earlier to compute the logarithm under the family of metrics, introduced by Hüper et al. [KIS21], which includes both the Canonical and the Euclidean metric.

Defining the Grassmann manifold $\mathrm{Gr}(n, k)$ as a quotient on $\mathrm{St}(n, k)$, it does not matter if we use the canonical or Euclidean metric. The induced metric on $\mathrm{Gr}(n, k)$ is the same.

The *Segre manifold*, $\text{Seg}(\mathbb{R}^{N_1} \times \cdots \times \mathbb{R}^{N_m})$, is the space of rank-1 tensors in the product space $\mathbb{R}^{N_1} \otimes \cdots \otimes \mathbb{R}^{N_m}$ [Har92]. It is immersed by

$$\begin{aligned} \text{Seg}: \mathbb{R}^+ \times S^{N_1-1} \times \cdots \times S^{N_m-1} &\rightarrow \mathbb{R}^{N_1} \otimes \cdots \otimes \mathbb{R}^{N_m}, \\ (\lambda, x_1, \dots, x_m) &\mapsto \lambda x_1 \otimes \cdots \otimes x_m. \end{aligned} \quad (5.76)$$

Swijsen, Van der Veken, and Vannieuwenhoven [SvV22] derive the exponential map for this manifold in the metric induced by the immersion. Swijsen [Swi22, theorem 6.2.1] derives the manifold logarithm. Both the exponential and logarithmic map are shown to be efficient to evaluate. Moreover, Swijsen [Swi22, corollary 6.2.4] shows that the sectional curvature satisfies

$$-\frac{1}{\lambda^2} \leq K \leq 0. \quad (5.77)$$

Note that in this case there is only a local lower bound, namely if the image of f fits in a chart with λ lower bounded by λ_* , then we may choose $H = -1/\lambda_*^2$.

The set of fixed-rank matrices is an embedded submanifold of $\mathbb{R}^{n \times n}$. Feppon and Lermusiaux [FL18, theorem 24] derive its sectional curvature. Similarly to the Segre manifold, the curvature is badly behaved close to the origin. Its geodesics are not known.

For embedded submanifolds, we also have some tools to construct retractions. Most immediately, projections³ of retractions on the supermanifold are retractions. For example, on the sphere $S^n \subset \mathbb{R}^{n+1}$,

$$\begin{aligned} R_x: T_x S^n &\rightarrow S^n, \\ y &\mapsto \frac{x+y}{\|x+y\|} \end{aligned} \quad (5.78)$$

is a retraction obtained from the (trivial) Riemannian exponential on \mathbb{R}^{n+1} . Similarly, the Eckart–Young theorem implies that the map

$$R_A: T_A \text{rank}^{-1}(r) \rightarrow \text{rank}^{-1}(r), \quad (5.79)$$

$$Y \mapsto \text{tsvd}(A+Y), \quad (5.80)$$

is a retraction, where tsvd is the r -truncated SVD. Absil and Malick [AM12] discuss a general framework for projection-based retractions. Absil and Oseledets [AO15] survey different retractions on the fixed-rank matrix manifold.

³Here, by *projection*, we mean the closest point on the submanifold. This is well-defined in a neighbourhood of the submanifold.

5.4.4 Warped products and Alexandrov spaces

The last class of manifolds that we mention is *warped products*. They are smooth product manifolds $M \times N$ such that the metric is of the form

$$\langle (X, Y), (X', Y') \rangle_{(p,q)} = \langle X, X' \rangle_p + f(p)^2 \langle Y, Y' \rangle_q \tag{5.81}$$

for some function $f: M \rightarrow \mathbb{R}$.

For warped products, there are particularly simple expressions for the curvature and geodesic equation. See for example [ONe83, propositions 7.38 and 7.42]. By identifying the Segre manifold as a warped product of \mathbb{R}^+ and $S^{N_1} \times \dots \times S^{N_m-1}$ with warping factor $f(\lambda) = \lambda$, Swijsen et al. [SvV22] are able to solve for the geodesics. Warped products where the first factor is \mathbb{R} are known as *cones*, and cones with this specific warping function are known as *linear cones*.

By considering other values of the warping parameter, Jacobsson et al. [Jac+24] consider a family of Riemannian manifolds that generalize the Segre manifold. Similarly to the Segre manifold, efficient expressions for the exponential and logarithmic map are derived. There is also a closed-form expression for the sectional curvature in terms of λ and the parameter [Jac+24, corollary 5.3]. From this expression it follows that the sectional curvature is lower bounded by

$$-\frac{1}{\lambda^2} \leq K. \tag{5.82}$$

A more theoretical approach to both upper and lower bound the curvature in warped products can be found in for example Alexander and Bishop [AB04]. Their theorem 1.1 for example implies the upper bound in (5.77). Their work comes from the study of *Alexandrov spaces*, which are generalizations of Riemannian manifolds with bounded sectional curvature. Formally, an Alexandrov space (often denoted $CBB(K)$) is a length space where geodesics satisfy the conclusion of Toponogov’s theorem with curvature bound K . Similarly, a $CAT(K)$ space is a length space where geodesics satisfy the conclusion of the companion theorem proposition 2.12 with curvature bound K . These spaces are more general than Riemannian manifolds and their intrinsic dimension can vary from point to point and even be infinite. Alexander, Kapovitch, and Petrunin [AKP23] is an introduction to the subject.

5.5 A concrete algorithm

We now want to approximate a function $f: [-1, 1]^m \rightarrow M$, where M is a Riemannian manifold. In the notation of section 5.3, $R = [-1, 1]^m$. algorithm 1 details a concrete realization of the approximation template discussed in section 5.1 that uses tensorized Chebyshev interpolation. Step 1 in the template corresponds to line 1 of algorithm 1, step 2 corresponds to lines 2–17, and step 3 corresponds to lines 18–19. We now describe steps 1 and 2 detail. The final step consists of applying the exponential map to the constructed approximation.

Algorithm 1 Approximating a map into a manifold.

Require: $f: [-1, 1]^m \rightarrow M$, number $N_k + 1$ of Chebyshev nodes to use in each direction $k = 1, \dots, m$.

Ensure: Approximant $\hat{f}: [-1, 1]^m \rightarrow M$ of f satisfies the error bound of theorem 5.5.

- 1: Choose $p \in M$ as (an approximation of) the Karcher mean of a large number of $f(x)$ for x chosen from a uniform distribution on $[-1, 1]^m$.
 - 2: Let $g = \log_p \circ f$.
 - 3: **for** $k = 1, \dots, m$ **do**
 - 4: **for** $i_k = 1, \dots, N_k + 1$ **do**
 - 5: Let $t[i_k] = \cos((2i_k - 1)\pi/(2(N_k + 1)))$ be the i_k th Chebyshev node.
 - 6: **end for**
 - 7: **end for**
 - 8: **for** $j = 1, \dots, n$ **do**
 - 9: Let $G_{i_1 \dots i_m j} = g_j(t[i_1], \dots, t[i_m])$.
 - 10: **end for**
 - 11: Define the core tensor C and the matrices $U^{(1)}, \dots, U^{(m)}, V$ by letting $\hat{G}_{i_1 \dots i_m j} = \sum_{a_1, \dots, a_m, b} C_{a_1 \dots a_m b} U_{i_1 a_1}^{(1)} \dots U_{i_m a_m}^{(m)} V_{b j}$ be an ST-HOSVD of G .
 - 12: **for** $k = 1, \dots, m$ **do**
 - 13: Let $h_{a_k}^{(k)}$ be the degree- N_k interpolating polynomial satisfying $h_{a_k}^{(k)}(t_{i_k}) = U_{i_k a_k}^{(k)}$.
 - 14: **end for**
 - 15: **for** $j = 1, \dots, n$ **do**
 - 16: Let $\hat{g}_j(x) = \sum_{a_1, \dots, a_m, b} C_{a_1 \dots a_m b} h_{a_1}^{(1)}(x_1) \dots h_{a_m}^{(m)}(x_m) V_{b j}$.
 - 17: **end for**
 - 18: Let $\hat{f} = \exp_p \circ \hat{g}$.
 - 19: Return \hat{f} .
-

5.5.1 Choosing from where to linearize

In practice, the choice of the point $p \in M$ in theorem 5.5 affects the error a lot. proposition 5.6 says that the further away from p we get, the bigger the error can be. It is thus natural to minimize the distance to p . Selecting the optimal p is a non-trivial problem, and is out of scope for this article. However, a related problem that has been studied in detail is the generalization of means and medians to manifolds. We will leverage this to define a heuristic for choosing p .

Consider a sequence $(\xi_i)_{i=1}^N$ of points in R , sampled independently from the same distribution on R . Then, consider the total squared distance to p ,

$$\sum_{i=1}^N d_M(f(\xi_i), p)^2. \quad (5.83)$$

A p that minimizes (5.83) is called a *Karcher mean*, or *Riemannian center of mass*, of $(f(\xi_i))_{i=1}^N$ [Kar77]. There are many situations where it is guaranteed to be unique, given the data ξ_i . For example, if there is a point in M whose distances to the ξ_i are all less than half the injectivity radius of the exponential map.⁴ In practice, the Karcher mean can be found via gradient descent [NB13, section 7.3.4], but also approximated efficiently using successive geodesic interpolation [CV15; CV19; CV20; Che+16], namely

$$m(p_1) = p_1, \quad (5.84)$$

$$m(p_1, \dots, p_N) = \exp_{p_N} \left(\frac{N-1}{N} \log_{p_N} m(p_1, \dots, p_{N-1}) \right). \quad (5.85)$$

For a general Riemannian manifold M , $m(p_1, \dots, p_N)$ depends on the order of p_1, \dots, p_N , but note that in the Euclidean case it is just the arithmetic mean of $(p_i)_{i=1}^N$, and thus exactly the Karcher mean.

In our implementation, we choose p as $m(f(\xi_1), \dots, f(\xi_N))$ with $(\xi_i)_{i=1}^N$ sampled from a uniform distribution on R .

5.5.2 Tensorized Chebyshev interpolation

For the second step in our implementation, we choose the *tensorized Chebyshev interpolation* discussed in section 5.2.2.

⁴This can be shown by combining Klingenberg's lemma [Pet16, lemma 6.4.7] with Kendall's uniqueness result [ABY13, (8.5)].

We have chosen to implement algorithm 1 with a basic HOSVD [DDV00a], computed with the sequentially truncated higher-order singular value decomposition (ST-HOSVD) approximation algorithm [VVM12] from the full tensor.

We mention also that the scheme is naturally extended to approximate vector-valued g by applying it component-wise.

5.6 Application: Krylov subspaces

In this example we approximate a map into the *Grassmann manifold* $\text{Gr}(n, k)$ of k -dimensional linear subspaces of \mathbb{R}^n . The example is inspired by Zimmermann's [ZH22, section 4.5] discussion of Grassmann manifolds in the context of model order reduction, and Benner, Gugercin, and Willcox's [BGW15] discussion in the context of Krylov subspace methods. See also Soodhalter, de Sturler, and Kilmer's [SSK20] recent survey on *Krylov subspace recycling*.

Consider the parametrized differential equation on $[0, 1]$,

$$\frac{d^2y}{dt^2} + x_1 \frac{dy}{dt} + x_2 y = t(1-t), \quad y(0) = y(1) = 0, \quad (5.86)$$

where x_1 and x_2 are parameters. This can model, for example, heat transfer in an uninsulated rod. We can discretize the equation by dividing $[0, 1]$ into n equal parts. Denoting $\Delta = 1/n$, $Y = (y(i\Delta))_{i=0}^n$, and $B = ((i\Delta)(1-i\Delta))_{i=0}^n$ and using a finite difference scheme to estimate d^2y/dt^2 and dy/dt , we have

$$\frac{Y_{i-1} - 2Y_i + Y_{i+1}}{\Delta^2} + x_1 \frac{-Y_{i-1} + Y_{i+1}}{2\Delta} + x_2 Y_i = B_i, \quad (5.87)$$

$$A(x_1, x_2)Y = B, \quad (5.88)$$

where

$$\begin{aligned} & A(x_1, x_2) \\ &= \frac{1}{\Delta^2} \begin{bmatrix} -2 & 1 & & & & & \\ 1 & -2 & 1 & & & & \\ & & \ddots & \ddots & \ddots & & \\ & & & 1 & -2 & 1 & \\ & & & & & 1 & -2 \end{bmatrix} + \frac{x_1}{2\Delta} \begin{bmatrix} 0 & 1 & & & & & \\ -1 & 0 & 1 & & & & \\ & & \ddots & \ddots & \ddots & & \\ & & & -1 & 0 & 1 & \\ & & & & & -1 & 0 \end{bmatrix} + x_2 \begin{bmatrix} 1 & & & & & & \\ & \ddots & & & & & \\ & & \ddots & & & & \\ & & & \ddots & & & \\ & & & & \ddots & & \\ & & & & & \ddots & \\ & & & & & & 1 \end{bmatrix}. \quad (5.89) \end{aligned}$$

Krylov subspace methods solve high-dimensional linear equations like these by looking for a solution in a *Krylov subspace* $\text{span}\{v, Av, \dots, A^{k-1}v\}$. Typically, A is also multiplied by some *preconditioner* P^{-1} , so that the Krylov

subspace is

$$\mathcal{K}_k(A, v) = \text{span} \{ v, P^{-1}Av, \dots, (P^{-1}A)^{k-1}v \}. \quad (5.90)$$

The tighter that the eigenvalues of $P^{-1}A$ are clustered, the faster the error in the Krylov subspace converges. In our numerical experiments, we use a generic algebraic multigrid preconditioner from Julia package `Preconditioners.jl`, but problem-specific preconditioners are available for most classic differential equations.

If we fix k and $v \in \mathbb{R}^n$, then

$$f(x_1, x_2) = \mathcal{K}_k(A(x_1, x_2), v) \quad (5.91)$$

is a map from \mathbb{R}^2 into $\text{Gr}(n, k)$. Rather than computing the Krylov subspace from its definition for each (x_1, x_2) , we want to be able to quickly compute an approximate Krylov subspace.




Figure 5.5: Approximation error for the map in (5.91) into the Grassmann manifold $\text{Gr}(n, k)$, compared against what is predicted by theorem 5.5, for different numbers of Chebyshev nodes N .

In this example, we approximate f on $R = [1, 2]^2$ with $n = 1000$, $k = 5$, and $v = B$. These parameters are chosen so that $AY = B$ can be solved with forward error on the order of 10^{-7} using $\mathcal{K}_k(A, v)$, and so that the original differential equation, (5.86), can be solved with error on the order of 10^{-5} .

We now proceed as in algorithm 1. Choose p and let $g = \log_p \circ f$. Then consider the polynomial \hat{g} that interpolates g in $(N + 1)^2$ Chebyshev nodes on R . Finally, let $\hat{f} = \exp_p \circ \hat{g}$.

Since both f and the manifold logarithm are analytic, and because of proposition 5.4, we expect the error $\|g - \hat{g}\|_p$ to decay at least exponentially in

N . Indeed, the numerical experiment presented in figure 5.5 supports this guess. In this experiment, we restrict f to a discrete subset $R' \subset R$ of 1000 points, sampled independently from a uniform distribution on R . We are thus able to use $\epsilon = \max_{x \in R'} \|g(x) - \widehat{g}(x)\|_p$ in (5.28) to derive a bound for $d(f(x), \widehat{f}(x))$. figure 5.5 shows that this error bound is confirmed by the experiment. After $N = 12$, rounding causes the errors to plateau and the error bound stops being meaningful. The code to produce figure 5.5 is available as `Example3.jl` in `ManiFactor.jl`.

5.7 Application: dynamical low-rank approximation

Expanding on a synthetic example by Ceruti and Lubich [CC22], we consider

$$A(x_1, x_2, x_3) = e^{x_1} e^{x_2 W_1} D(e^{x_3 W_2})^\top, \quad (5.92)$$

where D is an $n \times n$ diagonal matrix with diagonal entries $1, 2^{-1}, \dots, 2^{1-n}$, and W_1 and W_2 are skew-symmetric matrices, with entries sampled independently from a uniform distribution on $[-1, 1]$, and then normalized. Recall that when W is skew-symmetric, e^W is orthogonal. The best rank-1 approximation to A is

$$f(x) = e^{x_1} e^{x_2 W_1} e_1 (e^{x_3 W_2} e_1)^\top, \quad (5.93)$$

where e_1 is the column vector $(1, 0, \dots, 0)^\top$. Thus f is a map from \mathbb{R}^3 into the *Segre manifold* of rank-1 matrices $\text{Seg}(\mathbb{R}^n \times \mathbb{R}^n)$. We approximate f on $R = [-1, 1]^3$ with $n = 100$.

Similarly to section 5.6, we proceed as algorithm 1, while restricting to a discrete subset $R' \subset R$ of 1000 points sampled independently from a uniform distribution on R . Again, because of proposition 5.4, we expect the error to decay exponentially in the number of Chebyshev nodes $N = N_1 = N_2 = N_3$. By (5.82), we have a lower bound $H = -e^2$ for the sectional curvature. Thus theorem 5.5 implies that the approximation error is bounded by (5.29).

The numerical experiment presented in figure 5.6 again confirms the error bound. After $N = 13$, rounding causes the errors to plateau and the error bound stops being meaningful. The code to produce figure 5.6 is available as `Example5.jl` in `ManiFactor.jl`.




Figure 5.6: Approximation error for the map in (5.93) into the Segre manifold $\text{Seg}(\mathbb{R}^n \times \mathbb{R}^n)$, compared against what is predicted by theorem 5.5, for different numbers of Chebyshev nodes N in each variable.

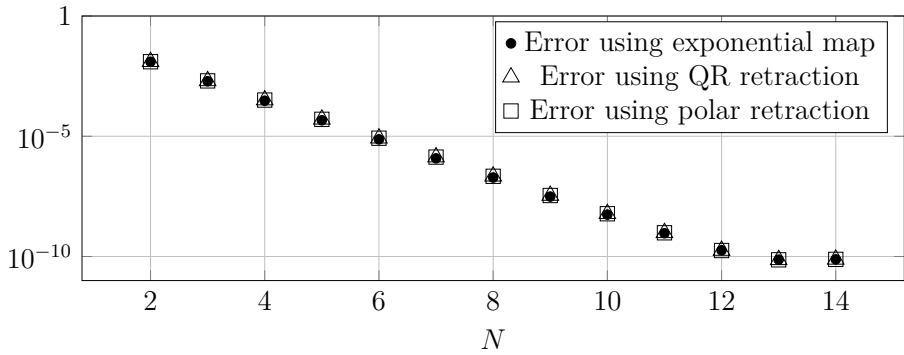


Figure 5.7: Approximation error for the map in (5.91), using different retraction methods.

5.8 Effects of using a retraction

In section 5.3.1, we discussed how a retraction and its inverse may be used instead of exponential and logarithmic maps. For manifolds where the exponential or logarithmic maps are not known, we have no choice but to retract. However, it can also be interesting to do so even when the geodesics are known, since it may be faster or more numerically stable.

`Manifolds.jl` implements two retractions on the Grassmann manifold, the *QR retraction* and the *polar retraction*. They are described in detail by Absil et al. [AMS08, section 4.1]. We are now interested in how replacing the exponential

map affects the Riemannian error.

Recall the setup from section 5.6. In figure 5.7, we recreate the experiment from figure 5.5 and compare with using the retractions. The discrepancy between the errors is almost imperceptible from the figure. Most importantly, they converge at the same rate. The largest discrepancy for either QR or polar is when $N = 7$. Then the approximation error when retracting is $1.47 \times 10^{-6} / (1.31 \times 10^{-6}) = 1.11$ times the error when using the exponential map. Since retractions, by definition, agree with the exponential map up to order 1, the discrepancy decreases as the error gets smaller. When $N = 14$, it is $9.568 \times 10^{-11} / (9.564 \times 10^{-11}) = 1.0004$. We emphasize that in the context of approximating a function to the 10th digit, this discrepancy is tiny.




Figure 5.8: Time it takes to evaluate $\hat{f}(x)$, using different retraction methods.

We are also interested in how the time to evaluate \hat{f} is affected. figure 5.8 shows the time it takes to evaluate \hat{f} in figure 5.7. It is measured using the `@benchmark` method from package `BenchmarkTools.jl`. The time to evaluate \hat{f} is the time it takes to evaluate \hat{g} plus the time it takes to evaluate the retraction, where the latter is independent of N . Thus retracting is significantly faster when N is small and \hat{g} is fast to evaluate. However, even though retracting is always faster, the speedup is marginal when compared to the time it takes to evaluate \hat{g} for large N .

The code to produce figures 5.7 and 5.8 is available as `Example4.jl` i `ManiFactor.jl`.

5.9 Conclusion and outlook

In conclusion, we find that the approximability of maps into a Riemannian manifold, when using the pullback-approximate-pushforward template, is related to the curvature of the manifold. Moreover, for many manifolds that are relevant for numerical analysis, especially matrix manifolds, explicit curvature bounds allow us to derive explicit error bounds.

We overviewed multivariate approximation theory, and recalled the mature theory about approximating real-valued functions. The theory of approximating maps into manifolds is basic in comparison. Before the current work, essentially only first-order error bounds were known. At the same time, the problem of approximating a map into a manifold is ubiquitous in numerical linear algebra. We argue that the error such an approximation introduces should be thought of intrinsically on the manifold, rather than in some ambient space or in some specific coordinates.

We also discussed some specific manifolds whose Riemannian exponential and logarithm are known, along with what is known about their curvature and convexity radii. From this discussion, we can conclude that many manifolds that are relevant in numerical linear algebra are “well-behaved” in the sense that our main theorem, theorem 5.5, is applicable and predicts a modest error.

We moreover presented two proof-of-concept applications: Krylov subspace recycling and dynamical low-rank approximation. In both of these applications, we had all the tools we needed to reason about the Riemannian errors, and the predicted error bound from theorem 5.5 was confirmed by numerical experiments.

From here, several areas of further research are visible. First, one could derive the geodesics and sectional curvature for more manifolds. For example spaces of tensors with fixed rank or multilinear rank. Second, one could develop more efficient and more optimal ways of choosing the point p from where the manifold is linearized. The current way, geodesic interpolation of sample evaluations, is a heuristic to estimate the Karcher mean of those evaluations, which in turn is a heuristic to minimize σ in theorem 5.5. One suggestion could be to approximate $\sigma \approx \log\left(\sum_{i=1}^N \exp(d_M(p, p_i))\right)$ and then optimize this with respect to p . The idea is that the suggested function is smooth in p , in contrast to the true σ , and so tools from Riemannian optimization can be used. Third, it could be possible to consider more than one tangent space. Given several approximations of f , based around several points, what is a good estimate of $f(x)$, and what are good ways of choosing such points? This research direction is

already started by Wang et al. [Wan+25]. They introduce the *multiple tangent space model*.

Recall from the start of this chapter that I am the main author of the work in this chapter, and did most of the research. The contributions of the other authors were mostly in technical discussions and feedback.

6. *Julia implementations*

`Manifolds.jl` [ABB24] is a popular software library for manifolds. It implements most of the manifolds that we have mentioned in this thesis, along with their exponential and logarithmic maps, distance functions, retractions, etc. We therefore chose to implement the Segre–Veronese and the other low-rank manifolds there.

The interface is very straightforward. To compute the exponential map of a tangent vector v at a point p on the valence- (n_1, n_2, n_3) Segre manifold, we just write

```
1 M = Segre((n1, n2, n3))
2 q = exp(M, p, v)
```

Internally, a point on the Segre is represented as a list of vectors. For example, if $n_1 = n_2 = n_3 = 3$, we have

```
1 > println(q)
2 4-element Vector{Vector{Float64}}:
3  [1.3070030321915433]
4  [-0.7109788094651939, 0.69966212597465, -0.07058358143427179]
5  [0.35060855782370387, -0.9153439147798758, -0.19803827119078293]
6  [-0.18531104885259742, -0.5754180036566499, -0.7965889380608708]
```

We remark that the last three vectors live on spheres, i.e. their norm is 1. One could argue that it would be smarter to only specify two numbers for each mode, so that we have to store less information. This is however not as simple to work with, and can suffer from ill-conditioning phenomena similar to gimbal lock.

In chapter 5, we presented a template for approximating maps into Riemannian manifolds. To perform our numerical experiments, we implemented this template as a Julia library, `ManiFactor.jl`. The idea is that, by being compatible with `Manifolds.jl`, all of the manifolds that are implemented there are available to our approximation scheme.

Recall the three-step template from section 5.1, and compare it to the body of our main function:

```

1 function approximate(m::Int, M::AbstractManifold, f::Function;
2   ...)
3   ...
4   g = chart . f
5   ghat = base_approximate(m, manifold_dimension(M), g)
6   fhat = chart_inv . ghat
7   return fhat
8 end

```

Here, m is the dimension of the domain, M is the target manifold, and f is the function from $\mathbb{R}^m \rightarrow M$ to be approximated. `base_approximate` is just a method to approximate maps between linear spaces. The default is the tensorized Chebyshev interpolation that we described in section 5.5, but you can also bring your own and just pass it as an argument to `approximate`.

The package is straightforward to use. Here is a script to generate the approximation that is shown in figure 5.1:

```

1 using ManiFactor
2 using Manifolds: Sphere, get_point, StereographicAtlas
3
4 n = 2
5 M = Sphere(n)
6
7 m = 2
8 f(x) = get_point(M, StereographicAtlas(), :south, [x[1]^2 - x
9   [2]^2, -x[1] * x[2]])
9 fhat = approximate(m, M, f)

```

6.1 Tests

An interesting problem, once a Riemannian manifold M is implemented, is writing tests that confirm you actually have the correct exponential and logarithmic maps. If M is embedded in \mathbb{R}^n , then many things are straightforward to test. For example, the acceleration of geodesics should be orthogonal to the tangent space. However, we want to take a more intrinsic approach that we can use for manifolds where we do not have a natural embedding.

Our goal is to write down a list of statements that together imply this, but that can also be tested numerically. Let $e: (p: M) \rightarrow T_p M \rightarrow M$ and $l: (p: M) \rightarrow M \rightarrow T_p M$ be candidates for the exponential and logarithmic

maps respectively. We claim that the following is sufficient if satisfied for all $p \in M$ and sufficiently small $X \in T_pM$:

1. $e(p, 0) = p$,
2. $\|\gamma'(t)\|_{\gamma(t)} = \|X\|_p$ for all $t \in [-1, 1]$, where $\gamma(t) = e(p, tX)$,
3. $\frac{d}{d\alpha} \big|_{\alpha=0} \ell(\gamma_\alpha) = 0$ and $\frac{d^2}{d\alpha^2} \big|_{\alpha=0} \ell(\gamma_\alpha) > 0$ for any one-parameter variation γ_α of γ , and where ℓ is the length functional defined in (2.8),
4. $l(e(p, X)) = X$.

In order to test these numerically, we assume we have some way of sampling points and tangent vectors randomly on M .

Item 1 is clearly necessary and straightforward to test.

Item 2 is the condition that e is a radial isometry. It also implies that its derivative is injective. It can be implemented by using a finite difference scheme to differentiate γ .

Item 3 is the condition that geodesics are local minima of the length functional. A variation of γ can be implemented using *Bézier curves*. The quadratic Bézier curve in T_pM from 0 to X with control point $\alpha Y \in T_pM$ is

$$b_\alpha(t) = (1-t)^2 0 + 2(1-t)t\alpha Y + t^2 X. \quad (6.1)$$

Note that $b_\alpha(0) = 0$ and $b_\alpha(1) = X$, and that b_0 is a straight line, so that

$$\gamma_\alpha = e(p, b_\alpha(t)) \quad (6.2)$$

is a variation of γ . Computing $\ell(\gamma_\alpha)$ can be done using finite differences to estimate γ'_α and some standard quadrature rule like the trapezoidal rule. Importantly, $\alpha \mapsto \ell(\gamma_\alpha)$ is just a map from \mathbb{R} to \mathbb{R} , and so its derivatives are similarly straightforward to estimate.

Item 4 is also straightforward to test, but care needs to be taken so that X is chosen within the injectivity radius.

More than one derivation error has been spotted in this way.

7. Conclusion

Recall the goals we set up in section 1.3:

- Explore the geometry of low-rank tensors, as well as apply this geometry to integrate low-rank ordinary differential equations.
- Analyze the error of approximation schemes for maps into manifolds, using methods that are intrinsic to those manifolds. We want to do this analysis in a sufficiently general setting so that it covers previous works on approximating maps into manifolds.

The second item, we considered in chapters 3 and 4. We considered the natural Lie group action of multiplying by an invertible matrix in each mode, or equivalently a change of basis in each mode. This action preserves the CP, TT, and Tucker ranks. In the case of CP rank, we further had to restrict to maximal Tucker rank for the action to be transitive. Similarly for TT rank. For the Tucker rank, we found that the action is transitive only for unbalanced ranks. We then used this action to derive a smooth homogeneous manifold structure on each of these sets.

The main application of the homogeneous structure was to derive a so-called *canonical metric*, and to show that the exponential map under this metric can be computed efficiently. The exponential map can already be evaluated fairly efficiently by just lifting it to the action group. However, we notice that the tangent spaces at our choice of representatives all have a certain rank structure. This structure is then exploited so that we can compute the exponential even more efficiently.

We emphasize that the maximal Tucker rank condition is an open condition, in the sense that for example a generic CP rank r tensor will have maximal Tucker rank. On the other hand, the restriction on the Tucker rank to be unbalanced seems much more limiting. While such rank-restrictions are maybe

not as useful for general-purpose low-rank techniques, since they single out one mode in particular, we mention that they have one very natural interpretation: time series. In applications, the time dimension often dominates. For example, the *Cars* movie I have lying around in my `Downloads` directory is an element of $(\mathbb{R}^{167620} \otimes \mathbb{R}^{1280} \otimes \mathbb{R}^{528})^{\times 3}$, where 167 620 is the number of frames, 1 280 is width resolution, 528 is the height resolution, and 3 is the number of colors. Note that we do not claim that $167\,620 = 1\,280 \cdot 528$, but rather that using Tucker decomposition with rank (t_1, t_2, t_3) with $t_1 = t_2 t_3$ is not unreasonable.

We presented an example that CP decomposed a given tensor via optimization. In general, any algorithm that evolves on a fixed-CP-rank manifold that also preserves multilinear rank, and whose initial condition is generic, can be interpreted as an algorithm evolving on the homogeneous manifold Σ_r , and similarly for TT rank. One such example would be exponential integrators for dynamical low-rank approximation. There, you multiply by an invertible (in the range of the matrix exponential) matrix in each time step, and so the algorithm is manifestly multilinear-rank-preserving.

The first item, we considered in chapter 5. We settled on a class of manifolds with lower bounded sectional curvature and known exponential and logarithmic maps. In this setting, we were able to derive a Riemannian error bound, theorem 5.5, in terms of the error on the tangent space and a lower sectional curvature bound. Importantly, we show in proposition 5.6 that our error bound is tight in the sense that, on a model manifold, we can construct an error that almost achieves the bound. In particular, this implies that the condition number lives in a very narrow interval.

Theorem 5.5 and proposition 5.6 formalize the intuition that negative curvature is bad for conditioning due to geodesics spreading apart. In dynamical systems for example, negative curvature has been related to instability, and in geometrical integration, it can impose bounds on allowed step sizes.

We found that many of the common manifolds from applications fit into this framework. In particular, many manifolds like the Stiefel and Grassmann manifolds have non-negative curvature. In this case, the Riemannian errors are not larger than the tangent space error. We also noted that for some manifolds there are more than one Riemannian structure to choose from. This is exemplified by the difference between the warped product metrics and the canonical metrics that can be put on sets of rank-1 tensors. In the first case, there is no global lower bound for the sectional curvature and close to 0 it approaches negative infinity, while in the second case the sectional curvature is globally lower bounded by $-11/2$. However, in the first case the logarithmic

map is known, while in the second case it is not. In applications, a bit of judgement has to be used to make an appropriate choice.

For practical applications, it is sometimes important to be able to use an approximate exponential map, called a retraction. Our framework allows us to bound the error also in this situation in terms of the error between the retraction and the exponential. However, we don't expect this bound to be as sharp due to the way we prove it. This intuition is confirmed in section 5.8 where the QR and polar retractions are shown to perform very similarly to the exponential map in terms of approximation error.

While retractions are a important when the exponential or logarithmic maps are not known explicit, or when they are intractable to compute, we stress that in many situations they can be computed efficiently. Efficient geodesic computations is a central theme in chapter 3, where we use the low-rank structure of the tangent space to derive efficient expressions for the exponential map. Similarly, the exponential and logarithmic maps for the Segre–Veronese manifolds, derived in chapter 4, can be evaluated in a number of operations proportional to the dimension of the manifold. It is not possible for these maps to have lower asymptotic computational cost.

We emphasize that the template we presented at the beginning of chapter 5 is widely used in practice. Thus our analysis is relevant in a host of different applications. For example, in the Krylov subspace recycling literature, it is uncommon to think about error between the actual subspaces. What we were able to show was that this is actually feasible, provided that one can interpret a Krylov subspace algorithm in terms of a linear approximation scheme and a retraction. Another example is dynamical low-rank approximation, which is a very active research field. Many dynamical low-rank methods can be directly interpreted as approximating some map into the manifold of fixed-rank matrices or tensors.

7.1 Future work

Theorem 5.5 is a statement about the approximation error in a single chart around a prescribed point. We discussed briefly a heuristic method for choosing this point as the Karcher mean of a set of function evaluations. However, there is a lot left to understand about what constitutes a *good* choice. In particular, when several points are chosen, how should they be chosen relative to each other? This work is already being pursued by Wang, Vandebril, Van der Veken, and Vannieuwenhoven [Wan+25], with what they term the *multiple tangent space model*.

One application for low-rank tensor manifolds, which we touched upon in section 3.5, is Riemannian optimization on low-rank tensor manifolds. We considered the problem of finding an approximate low-rank decomposition of a given tensor. A fundamental difficulty in such problems is to invert Riemannian metric so that we can map gradients to steepest-descent directions. Luckily, we found this can be done efficiently. Moreover, the expression we derived is not limited to the specific application that we considered, but it can be useful for arbitrary Riemannian optimization problems on our low-rank tensor manifolds.

Another avenue for future work is the application of the manifold structure on low-rank tensors to dynamical low-rank approximation. Future work in this direction should look at what can be done to overcome ill-conditioning and degeneracy in the decomposition factors. In particular, the presence of the inverse operation in expressions like (3.24) should be reviewed. One could also look at using the retraction induced by the Lie exponential, which has roughly half the computational cost of the *Riemannian exponential*.

The last avenue that we'll mention is a connection between what is called *identifiability* in the tensor decomposition literature, and what is called *stabilizer subgroup* in Riemannian geometry. When we derived the quotient expression for our low-rank CP tensors, we used Kruskal's theorem to find the denominator. However, Kruskal's is only the most classical theorem in identifiability. For example, a simple modification to our setup could be: what is the stabilizer when the change-of-basis group acts on sets of fixed-CP-rank fixed-Tucker-rank tensors if the Tucker rank is not maximal? Kruskal's theorem does not immediately answer this because if the Tucker rank is not maximal then the factors in each term are not all linearly independent. We do not expect the answer to be as straightforward as the maximal-Tucker-rank case. Answers to such questions often boil down to whether the situation can be reduced to the matrix case or not.

Bibliography

- [AB04] S. B. Alexander and R. L. Bishop. “Curvature bounds for warped products of metric spaces”. In: *Geometric & Functional Analysis GFAA* 14 (6 Dec. 1, 2004), pp. 1143–1181. DOI: 10.1007/s00039-004-0487-2.
- [ABB24] S. Axen, M. Baran, and R. Bergmann. *Manifolds.jl*. Version v0.10.10. Dec. 2024. DOI: 10.5281/zenodo.14538142.
- [ABY13] M. Arnaudon, F. Barbaresco, and L. Yang. “Medians and Means in Riemannian Geometry: Existence, Uniqueness and Computation”. In: *Matrix Information Geometry*. Ed. by F. Nielsen and R. Bhatia. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 169–197. ISBN: 978-3-642-30232-9. DOI: 10.1007/978-3-642-30232-9_8.
- [ADK11] E. Acar, D. M. Dunlavy, and T. G. Kolda. “A scalable optimization approach for fitting canonical tensor decompositions”. In: *Journal of chemometrics* 25.2 (2011), pp. 67–86.
- [AF68] N. Ahmed and P. Fisher. “Study of algorithmic properties of Chebyshev coefficients”. In: *International Journal of Computer Mathematics* 2.1-4 (1968), pp. 307–317.
- [Afs11] B. Afsari. “Riemannian L^p center of mass: Existence, uniqueness, and convexity”. In: *Proc. Amer. Math. Soc.* 139.02 (Feb. 2011), pp. 655–655. ISSN: 0002-9939. DOI: 10.1090/s0002-9939-2010-10541-5.
- [AKP23] S. Alexander, V. Kapovitch, and A. Petrunin. *Alexandrov geometry: foundations*. 2023. arXiv: 1903.08539 [math.DG]. URL: <https://arxiv.org/abs/1903.08539>.
- [AM12] P.-A. Absil and J. Malick. “Projection-like Retractions on Matrix Manifolds”. In: *SIAM Journal on Optimization* 22.1 (2012), pp. 135–158. DOI: 10.1137/100802529.

- [AM25] P.-A. Absil and S. Maitaigne. “The Ultimate Upper Bound on the Injectivity Radius of the Stiefel Manifold”. In: *SIAM Journal on Matrix Analysis and Applications* 46.2 (2025), pp. 1145–1167. DOI: 10.1137/24M1644808.
- [AMS08] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton, NJ: Princeton University Press, 2008, pp. xvi+224. ISBN: 978-0-691-13298-3.
- [And+14] E. Andruchow, G. Larotonda, L. Recht, and A. Varela. “The left invariant metric in the general linear group”. In: *Journal of Geometry and Physics* 86 (2014), pp. 241–257. ISSN: 0393-0440. DOI: 10.1016/j.geomphys.2014.08.009.
- [AO15] P.-A. Absil and I. V. Oseledets. “Low-rank retractions: a survey and new results”. In: *Computational Optimization and Applications* 62 (1 Sept. 1, 2015), pp. 5–29. DOI: 10.1007/s10589-014-9714-4.
- [Arv03] A. Arvanitogeorgos. *An Introduction to Lie Groups and the Geometry of Homogeneous Spaces*. Student mathematical library. American Mathematical Society, 2003. ISBN: 9780821827789.
- [Axe23] S. Axen. *The injectivity radii of the unitary groups*. <https://sethaxen.com/blog/2023/02/the-injectivity-radii-of-the-unitary-groups/>. Accessed: 2025-08-28. 2023.
- [Bal24] E. Ballico. “Joins, Secant Varieties and Their Associated Grassmannians”. In: *Mathematics* 12.9 (Apr. 2024), p. 1274. ISSN: 2227-7390. DOI: 10.3390/math12091274.
- [BBV19] C. Beltrán, P. Breiding, and N. Vannieuwenhoven. “Pencil-based algorithms for tensor rank decomposition are not stable”. In: *SIAM J. Matrix Anal. Appl.* 40.2 (2019), pp. 739–773. DOI: 10.1137/18m1200531.
- [BC24] A. T. Blomenhofer and A. Casarotti. *Nondefectivity of invariant secant varieties*. 2024. arXiv: 2312.12335 [math.AG].
- [BDS24] H. Ben-Zion Vardi, N. Dyn, and N. Sharon. “Hermite subdivision schemes for manifold-valued Hermite data”. In: *Computer Aided Geometric Design* 111 (2024), p. 102342. ISSN: 0167-8396. DOI: 10.1016/j.cagd.2024.102342.
- [Beb00] M. Bebendorf. “Approximation of boundary element matrices”. In: *Numerische Mathematik* 86.4 (2000), pp. 565–589.
- [BEM16] D. Bigoni, A. P. Engsig-Karup, and Y. M. Marzouk. “Spectral Tensor-Train Decomposition”. In: *SIAM Journal on Scientific Computing* 38.4 (Jan. 2016), A2405–A2439. ISSN: 1095-7197. DOI: 10.1137/15m1036919.

- [Ber+18] A. Bernardi, E. Carlini, M. V. Catalisano, A. Gimigliano, and A. Oneto. “The Hitchhiker guide to: Secant Varieties and Tensor Decomposition”. In: *Mathematics* 6.314 (2018). DOI: 10.3390/math6120314.
- [BG18] R. Bergmann and P.-Y. Gousenbourger. “A Variational Model for Data Fitting on Manifolds by Minimizing the Acceleration of a Bézier Curve”. In: *Frontiers in Applied Mathematics and Statistics* 4 (2018). ISSN: 2297-4687. DOI: 10.3389/fams.2018.00059.
- [BGM09] G. Beylkin, J. Garcke, and M. J. Mohlenkamp. “Multivariate Regression and Machine Learning with Sums of Separable Functions”. In: *SIAM Journal on Scientific Computing* 31.3 (2009), pp. 1840–1857.
- [BGW15] P. Benner, S. Gugercin, and K. Willcox. “A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems”. In: *SIAM Review* 57.4 (2015), pp. 483–531. DOI: 10.1137/130932715.
- [Bha07] R. Bhatia. *Positive Definite Matrices*. Princeton University Press, 2007. ISBN: 9780691129181. URL: <http://www.jstor.org/stable/j.ctt7rxv2> (visited on 07/03/2024).
- [Bla+09] S. Blanes, F. Casas, J. Oteo, and J. Ros. “The Magnus expansion and some of its applications”. In: *Physics Reports* 470.5 (2009), pp. 151–238. ISSN: 0370-1573. DOI: <https://doi.org/10.1016/j.physrep.2008.11.001>.
- [BM02] G. Beylkin and M. J. Mohlenkamp. “Numerical operator calculus in higher dimensions”. In: *Proceedings of the National Academy of Sciences of the United States of America* 99.16 (2002), pp. 10246–10251.
- [BM05] G. Beylkin and M. J. Mohlenkamp. “Algorithms for numerical analysis in high dimensions”. In: *SIAM Journal on Scientific Computing* 26.6 (2005), pp. 2133–2159.
- [Bou23] N. Boumal. *An introduction to optimization on smooth manifolds*. Cambridge University Press, 2023. DOI: 10.1017/9781009166164.
- [BP23] B. Buet and X. Pennec. *Flagfolds*. 2023. arXiv: 2305.10583 [math.CA].
- [Bre+23] T. G. K. Brett W. Bader et al. *Tensor Toolbox for MATLAB, Version 3.6*. 2023. URL: www.tensortoolbox.org.
- [BT04] J.-P. Berrut and L. N. Trefethen. “Barycentric Lagrange Interpolation”. In: *SIAM Review* 46.3 (2004), pp. 501–517. DOI: 10.1137/S0036144502417715.

- [BV18] P. Breiding and N. Vannieuwenhoven. “A Riemannian Trust Region Method for the Canonical Tensor Rank Approximation Problem”. In: *SIAM Journal on Optimization* 28.3 (2018), pp. 2435–2465. DOI: 10.1137/17M114618X.
- [BZ21] T. Bendokat and R. Zimmermann. *The real symplectic Stiefel and Grassmann manifolds: metrics, geodesics and applications*. 2021. arXiv: 2108.12447 [math.DG].
- [BZA24] T. Bendokat, R. Zimmermann, and P.-A. Absil. “A Grassmann manifold handbook: basic geometry and computational aspects”. In: *Advances in Computational Mathematics* 50 (1 Jan. 5, 2024). DOI: 10.1007/s10444-023-10090-8.
- [CC04] O. Christensen and K. L. Christensen. *Approximation Theory: from Taylor Polynomials to Wavelets*. Applied and Numerical Harmonic Analysis. Boston, MA: Birkhäuser, 2004. ISBN: 0-8176-3600-5.
- [CC22] G. Ceruti and L. Christian. “An unconventional robust integrator for dynamical low-rank approximation”. In: *BIT Numerical Mathematics* 62.1 (2022). DOI: 10.1007/s10543-021-00873-0.
- [CE08] J. Cheeger and D. G. Ebin. *Comparison Theorems in Riemannian Geometry*. English (US). AMS Chelsea Publishing, 2008. ISBN: 9780821844175.
- [Che+16] G. Cheng, J. Ho, H. Salehian, and B. C. Vemuri. “Recursive Computation of the Fréchet Mean on Non-positively Curved Riemannian Manifolds with Applications”. In: *Riemannian Computing in Computer Vision*. Ed. by P. K. Turaga and A. Srivastava. Cham: Springer International Publishing, 2016, pp. 21–43. ISBN: 978-3-319-22957-7. DOI: 10.1007/978-3-319-22957-7_2.
- [Che99] C.-H. Chen. “Warped products of metric spaces of curvature bounded from above”. In: *Transactions of the American Mathematical Society* 351.12 (Aug. 1999), pp. 4727–4740. ISSN: 1088-6850. DOI: 10.1090/s0002-9947-99-02154-6.
- [Cic+15] A. Cichocki, D. Mandic, L. De Lathauwer, G. Zhou, Q. Zhao, C. Caiafa, and H. A. PHAN. “Tensor Decompositions for Signal Processing Applications: From two-way to multiway component analysis”. In: *IEEE Signal Processing Magazine* 32.2 (2015), pp. 145–163. DOI: 10.1109/MSP.2013.2297439.
- [CJ10] P. Comon and C. Jutten. *Handbook of Blind Source Separation: Independent component analysis and applications*. Academic press, 2010.

- [CK25] A. Cortinovis and D. Kressner. *Adaptive randomized pivoting for column subset selection, DEIM, and low-rank approximation*. 2025. arXiv: 2412.13992 [math.NA].
- [CKL22] G. Ceruti, J. Kusch, and C. Lubich. “A rank-adaptive robust integrator for dynamical low-rank approximation”. In: *BIT Numerical Mathematics* 62.4 (2022), pp. 1149–1174.
- [CL10] J. R. Cardoso and F. S. Leite. “Exponentials of skew-symmetric matrices and logarithms of orthogonal matrices”. In: *Journal of computational and applied mathematics* 233.11 (2010), pp. 2867–2875.
- [CLD09] P. Comon, X. Luciani, and A. L. De Almeida. “Tensor decompositions, alternating least squares and other tales”. In: *Journal of Chemometrics: A Journal of the Chemometrics Society* 23.7-8 (2009), pp. 393–405.
- [CMM18] A. Casarotti, A. Massarenti, and M. Mella. “On Comon’s and Strassen’s Conjectures”. In: *Mathematics* 6.11 (2018). ISSN: 2227-7390. DOI: 10.3390/math6110217.
- [CMM19] C. Curry, S. Marsland, and R. I. McLachlan. “Principal symmetric space analysis”. In: *Journal of Computational Dynamics* 6.2 (2019), pp. 251–276. ISSN: 2158-2491. DOI: 10.3934/jcd.2019013.
- [CMO11] S. H. Christiansen, H. Z. Munthe-Kaas, and B. Owren. “Topics in structure-preserving discretization”. In: *Acta Numerica* 20 (2011), pp. 1–119. DOI: 10.1017/S096249291100002X.
- [CO03] E. Celledoni and B. Owren. “On the Implementation of Lie Group Methods on the Stiefel Manifold”. In: *Numerical Algorithms* 32 (2 Apr. 1, 2003), pp. 163–183. DOI: 10.1023/A:1024079724094.
- [Com+08] P. Comon, G. Golub, L.-H. Lim, and B. Mourrain. “Symmetric Tensors and Symmetric Tensor Rank”. In: *SIAM Journal on Matrix Analysis and Applications* 30.3 (2008), pp. 1254–1279. DOI: 10.1137/060661569.
- [Cor+16] E. Cornea, H. Zhu, P. Kim, and J. G. Ibrahim. “Regression Models on Riemannian Symmetric Spaces”. In: *Journal of the Royal Statistical Society Series B: Statistical Methodology* 79.2 (Mar. 2016), pp. 463–482. ISSN: 1369-7412. DOI: 10.1111/rssb.12169.
- [COV14] L. Chiantini, G. Ottaviani, and N. Vannieuwenhoven. “An Algorithm For Generic and Low-Rank Specific Identifiability of Complex Tensors”. In: *SIAM Journal on Matrix Analysis and Applications* 35.4 (2014), pp. 1265–1287. DOI: 10.1137/140961389.

- [CST23] D. Chafamo, V. Shanmugam, and N. Tokcan. “Robust Bayesian Tensor Factorization with Zero-Inflated Poisson Model and Consensus Aggregation”. In: *arXiv:2308.08060v1* (Aug. 15, 2023).
- [CSV16] R. Chakraborty, D. Seo, and B. C. Vemuri. “An Efficient Exact-PGA Algorithm for Constant Curvature Manifolds”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2016.
- [CV15] R. Chakraborty and B. C. Vemuri. “Recursive Fréchet Mean Computation on the Grassmannian and Its Applications to Computer Vision”. In: *2015 IEEE International Conference on Computer Vision (ICCV)*. 2015, pp. 4229–4237. DOI: 10.1109/ICCV.2015.481.
- [CV19] R. Chakraborty and B. C. Vemuri. “Statistics on the Stiefel manifold: Theory and applications”. In: *The Annals of Statistics* 47.1 (2019), pp. 415–438. DOI: 10.1214/18-AOS1692.
- [CV20] R. Chakraborty and B. C. Vemuri. “Efficient recursive estimation of the Riemannian barycenter on the hypersphere and the special orthogonal group with applications”. In: *Riemannian Geometric Statistics in Medical Image Analysis*. Ed. by X. Pennec, S. Sommer, and T. Fletcher. Cambridge, MA: Academic Press, 2020, pp. 273–297. ISBN: 978-0-12-814725-2. DOI: 10.1016/B978-0-12-814725-2.00015-7.
- [DCN25] W. Diepeveen, J. Chew, and D. Needell. “Curvature-corrected tangent space-based approximation of manifold-valued data”. In: *Information and Inference: A Journal of the IMA* 14.4 (2025), iaaf031.
- [DDV00a] L. De Lathauwer, B. De Moor, and J. Vandewalle. “A Multilinear Singular Value Decomposition”. In: *SIAM Journal on Matrix Analysis and Applications* 21 (2000), pp. 1253–1278.
- [DDV00b] L. De Lathauwer, B. De Moor, and J. Vandewalle. “On the Best Rank-1 and Rank-(R_1, R_2, \dots, R_N) Approximation of Higher-Order Tensors”. In: *SIAM Journal on Matrix Analysis and Applications* 21.4 (2000), pp. 1324–1342. DOI: 10.1137/S0895479898346995.
- [De +22] T. De Weer, N. Vannieuwenhoven, N. Lammens, and K. Meerbergen. “The parametrized superelement approach for lattice joint modelling and simulation”. In: *Computational Mechanics* 70.2 (May 2022), pp. 451–475. ISSN: 1432-0924. DOI: 10.1007/s00466-022-02176-9.

- [Den+25] Z. Deng, P.-A. Absil, K. A. Gallivan, and W. Huang. *The Exponential of Skew-Symmetric Matrices: A Nearby Inverse and Efficient Computation of Derivatives*. 2025. arXiv: 2506.18302 [math.DG].
- [Dib17] J. Dibble. “The convexity radius of a Riemannian manifold”. In: *Asian Journal of Mathematics* 21.1 (2017), pp. 169–174. DOI: 10.4310/AJM.2017.v21.n1.a4.
- [Die24] W. Diepeveen. *Pulling back symmetric Riemannian geometry for data analysis*. 2024. arXiv: 2403.06612 [math.DG]. URL: <https://arxiv.org/abs/2403.06612>.
- [DKS21] S. Dolgov, D. Kressner, and C. Strössner. “Functional Tucker Approximation Using Chebyshev Interpolation”. In: *SIAM Journal on Scientific Computing* 43.3 (2021), A2190–A2210. DOI: 10.1137/20M1356944.
- [Dra24] J. Draisma. “Erratum: A Counterexample to Comon’s Conjecture”. In: *SIAM Journal on Applied Algebra and Geometry* 8.1 (2024), pp. 225–225.
- [DS17a] N. Dyn and N. Sharon. “A global approach to the refinement of manifold data”. In: *Mathematics of Computation* 86 (2017). DOI: 10.1090/mcom/3087.
- [DS17b] N. Dyn and N. Sharon. “Manifold-valued subdivision schemes based on geodesic inductive averaging”. In: *Journal of Computational and Applied Mathematics* 311 (2017), pp. 54–67. ISSN: 0377-0427. DOI: 10.1016/j.cam.2016.07.008.
- [EAS98] A. Edelman, T. A. Arias, and S. T. Smith. “The geometry of algorithms with orthogonality constraints”. In: *SIAM journal on Matrix Analysis and Applications* 20.2 (1998), pp. 303–353.
- [ES09] L. Eldén and B. Savas. “A Newton–Grassmann Method for Computing the Best Multilinear Rank- (r_1, r_2, r_3) Approximation of a Tensor”. In: *SIAM Journal on Matrix Analysis and Applications* 31.2 (2009), pp. 248–271. DOI: 10.1137/070688316.
- [FL18] F. Feppon and P. F. J. Lermusiaux. “A Geometric Approach to Dynamical Model Order Reduction”. In: *SIAM Journal on Matrix Analysis and Applications* 39.1 (Jan. 2018), pp. 510–538. ISSN: 1095-7162. DOI: 10.1137/16m1095202.

- [FO16] E. Frolov and I. Oseledets. “Fifty Shades of Ratings: How to Benefit from a Negative Feedback in Top-N Recommendations Tasks”. In: *Proceedings of the 10th ACM Conference on Recommender Systems. RecSys '16*. Boston, Massachusetts, USA: Association for Computing Machinery, 2016, pp. 91–98. ISBN: 9781450340359. DOI: 10.1145/2959100.2959170.
- [FO17] E. Frolov and I. Oseledets. “Tensor methods and recommender systems”. In: *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 7.3 (2017), e1201.
- [FO23] E. Frolov and I. Oseledets. “Tensor-based sequential learning via Hankel matrix representation for next item recommendations”. In: *IEEE Access* 11 (2023), pp. 6357–6371.
- [GGT13] P. Gonnet, S. Guttel, and L. N. Trefethen. “Robust Padé approximation via SVD”. In: *SIAM Review* 55.1 (2013), pp. 101–117.
- [GH23] M. Griebel and H. Harbrecht. “Analysis of tensor approximation schemes for continuous functions”. In: *Foundations of Computational Mathematics* 23.1 (2023), pp. 219–240.
- [GKM19] A. Gorodetsky, S. Karaman, and Y. Marzouk. “A continuous analogue of the tensor-train decomposition”. In: *Computer Methods in Applied Mechanics and Engineering* 347 (2019), pp. 59–84. ISSN: 0045-7825. DOI: <https://doi.org/10.1016/j.cma.2018.12.015>.
- [GKT13] L. Grasedyck, D. Kressner, and C. Tobler. “A literature survey of low-rank tensor approximation techniques”. In: *GAMM-Mitteilungen* 36.1 (2013), pp. 53–78.
- [GL18] E. S. Gawlik and M. Leok. “Embedding-based interpolation on the special orthogonal group”. In: *SIAM Journal on Scientific Computing* 40.2 (2018), A721–A746.
- [GMA19] P.-Y. Gousenbourger, E. Massart, and P.-A. Absil. “Data Fitting on Manifolds with Composite Bézier-Like Curves and Blended Cubic Splines”. In: *Journal of Mathematical Imaging and Vision* 61 (2019), pp. 645–671. DOI: 10.1007/s10851-018-0865-2.
- [GOC25] M. Ghirardelli, B. Owren, and E. Celledoni. *Conditional Stability of the Euler Method on Riemannian Manifolds*. 2025. arXiv: 2503.09434 [math.NA].
- [Gor+10] S. A. Goreinov, I. V. Oseledets, D. V. Savostyanov, E. E. Tyrtyshnikov, and N. L. Zamarashkin. “How to find a good submatrix”. In: *Matrix Methods: Theory, Algorithms And Applications: Dedicated to the Memory of Gene Golub*. World Scientific, 2010, pp. 247–256.
- [Gre78] W. Greub. *Multilinear Algebra*. 2nd ed. Springer-Verlag, 1978. DOI: 10.1007/978-1-4613-9425-9.

- [GS+00] I. M. Gelfand, R. A. Silverman, et al. *Calculus of Variations*. Courier Corporation, 2000.
- [GS13] P. Grohs and M. Sprecher. “Projection-based quasiinterpolation in manifolds”. In: *SAM Report 23* (2013).
- [GSS23] C. G. Gebhardt, J. Schubert, and M. C. Steinbach. “Long-time principal geodesic analysis in director-based dynamics of hybrid mechanical systems”. In: *Communications in Nonlinear Science and Numerical Simulation* 122 (2023), p. 107240. ISSN: 1007-5704. DOI: 10.1016/j.cnsns.2023.107240.
- [GSY17] P. Grohs, M. Sprecher, and T. Yu. “Scattered manifold-valued data approximation”. In: *Numer. Math.* 135.4 (Apr. 2017), pp. 987–1010. ISSN: 0029-599X. DOI: 10.1007/s00211-016-0823-0.
- [GT01] S. A. Goreinov and E. E. Tyrtyshnikov. “The maximal-volume concept in approximation by low-rank matrices”. In: *Contemporary Mathematics* 280 (2001), pp. 47–52.
- [GTZ97] S. Goreinov, E. Tyrtyshnikov, and N. Zamarashkin. “A theory of pseudoskeleton approximations”. In: *Linear Algebra and its Applications* 261.1 (1997), pp. 1–21. ISSN: 0024-3795. DOI: 10.1016/S0024-3795(96)00301-1.
- [GX03] J. Gallier and D. Xu. “Computing exponentials of skew-symmetric matrices and logarithms of orthogonal matrices”. In: *International Journal of Robotics and Automation* 18.1 (2003), pp. 10–20.
- [Hac19] W. Hackbusch. *Tensor Spaces and Numerical Tensor Calculus*. Ed. by R. Bank, R. Graham, J. Stoer, R. Varga, and H. Yserentant. 2nd ed. Springer Series in Computational Mathematics 42. Springer-Verlag, 2019. DOI: 10.1007/978-3-030-35554-8.
- [Har92] J. Harris. *Algebraic Geometry, A First Course*. Vol. 133. Graduate Text in Mathematics. New York, NY: Springer-Verlag, 1992, p. 328.
- [Hås89] J. Håstad. “Tensor rank is NP-complete”. In: *International colloquium on automata, languages, and programming*. Springer. 1989, pp. 451–460.
- [HFJ14] J. Hinkle, P. T. Fletcher, and S. Joshi. “Intrinsic Polynomials for Regression on Riemannian Manifolds”. In: *Journal of Mathematical Imaging and Vision* 50 (1 2014), pp. 32–52. DOI: 10.1007/s10851-013-0489-5.
- [Hig08] N. J. Higham. *Functions of Matrices*. Society for Industrial and Applied Mathematics, 2008. DOI: 10.1137/1.9780898717778.

- [HKD20] D. Hong, T. G. Kolda, and J. A. Duersch. “Generalized Canonical Polyadic Tensor Decomposition”. In: *SIAM Review* 62.1 (2020), pp. 133–163. DOI: 10.1137/18M1203626.
- [HL13] C. J. Hillar and L.-H. Lim. “Most Tensor Problems Are NP-Hard”. In: *J. ACM* 60.6 (Nov. 2013). ISSN: 0004-5411. DOI: 10.1145/2512329.
- [HL18] S. Hu and G. Li. “Convergence rate analysis for the higher order power method in best rank one approximations of tensors”. In: *Numerische Mathematik* 140.4 (2018), pp. 993–1031.
- [HL23] R. Hielscher and L. Lippert. “Approximating the derivative of manifold-valued functions”. In: *Journal of Approximation Theory* 285 (Jan. 2023), p. 105832. ISSN: 0021-9045. DOI: 10.1016/j.jat.2022.105832.
- [HLS98] M. Hochbruck, C. Lubich, and H. Selhofer. “Exponential Integrators for Large Systems of Differential Equations”. In: *SIAM Journal on Scientific Computing* 19.5 (1998), pp. 1552–1574. DOI: 10.1137/S1064827595295337.
- [HMT11] N. Halko, P.-G. Martinsson, and J. A. Tropp. “Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions”. In: *SIAM Review* 53.2 (2011), pp. 217–288.
- [HN18] B. Hashemi and Y. Nakatsukasa. “On the spectral problem for trivariate functions”. In: *BIT Numerical Mathematics* 58 (2018), pp. 981–1008.
- [HNS23] M. Hochbruck, M. Neher, and S. Schrammer. “Dynamical low-rank integrators for second-order matrix differential equations”. In: *BIT* 63.1 (Jan. 2023). ISSN: 1572-9125. DOI: 10.1007/s10543-023-00941-7.
- [HRS12a] S. Holtz, T. Rohwedder, and R. Schneider. “On manifolds of tensors of fixed TT-rank”. In: *Numerische Mathematik* 120 (4 Apr. 1, 2012), pp. 701–731. DOI: 10.1007/s00211-011-0419-7.
- [HRS12b] S. Holtz, T. Rohwedder, and R. Schneider. “The Alternating Linear Scheme for Tensor Optimization in the Tensor Train Format”. In: *SIAM Journal on Scientific Computing* 34.2 (2012), A683–A713. DOI: 10.1137/100818893.
- [HRW19] B. Heeren, M. Rumpf, and B. Wirth. “Variational time discretization of Riemannian splines”. In: *IMA Journal of Numerical Analysis* 39.1 (2019), pp. 61–104. ISSN: 0272-4979. DOI: 10.1093/imanum/drx077.

- [HT17] B. Hashemi and L. N. Trefethen. “Chebfun in three dimensions”. In: *SIAM Journal on Scientific Computing* 39.5 (2017), pp. C341–C363.
- [Hu+25] N. Hu, M. A. Iwen, D. Needell, and R. Wang. *Convergence of the alternating least squares algorithm for CP tensor decompositions*. 2025. arXiv: 2505.14037 [math.NA].
- [HW21] H. Hardering and B. Wirth. “Quartic L^p -convergence of cubic Riemannian splines”. In: *IMA Journal of Numerical Analysis* 42.4 (Oct. 2021), pp. 3360–3385. ISSN: 0272-4979. DOI: 10.1093/imanum/drab077. eprint: <https://academic.oup.com/imanum/article-pdf/42/4/3360/46323898/drab077.pdf>.
- [Ise+00] A. Iserles, H. Z. Munthe-Kaas, S. P. Nørsett, and A. Zanna. “Lie-group methods”. In: *Acta Numerica* 9 (2000), pp. 215–365. DOI: 10.1017/S0962492900002154.
- [Ish+09] M. Ishteva, L. De Lathauwer, P.-A. Absil, and S. Van Huffel. “Differential-geometric Newton method for the best rank-(R 1, R 2, R 3) approximation of tensors”. In: *Numerical Algorithms* 51.2 (2009), pp. 179–194.
- [Ish+11] M. Ishteva, P.-A. Absil, S. Van Huffel, and L. De Lathauwer. “Best Low Multilinear Rank Approximation of Higher-Order Tensors, Based on the Riemannian Trust-Region Scheme”. In: *SIAM Journal on Matrix Analysis and Applications* 32.1 (2011), pp. 115–135. DOI: 10.1137/090764827.
- [Jac] S. Jacobsson. <https://gitlab.kuleuven.be/numa/public/alpha-warped-experiments>.
- [Jac+24] S. Jacobsson, L. Swijsen, J. V. der Veken, and N. Vannieuwenhoven. *Warped geometries of Segre-Veronese manifolds*. 2024. arXiv: 2410.00664 [math.NA].
- [Jac+25] S. Jacobsson, R. Vandebril, J. Van der Veken, and N. Vannieuwenhoven. “Approximating maps into manifolds with lower curvature bounds”. In: *BIT Numerical Mathematics* 65.3 (2025), p. 37.
- [Jac25] S. Jacobsson. *A homogeneous geometry of low-rank tensors*. 2025. arXiv: 2512.13594 [math.NA].
- [JZ25] R. Jensen and R. Zimmermann. *Maximum volume coordinates for Grassmann interpolation: Lagrange, Hermite, and errors*. 2025. arXiv: 2506.01574 [math.NA].
- [Kar77] H. Karcher. “Riemannian center of mass and mollifier smoothing”. In: *Communications on Pure and Applied Mathematics* 30.5 (Sept. 1977), pp. 509–541. ISSN: 1097-0312. DOI: 10.1002/cpa.3160300502.

- [Kar87] H. Karcher. *Riemannian comparison constructions*. SFB 256, 1987.
- [KB09] T. G. Kolda and B. W. Bader. “Tensor Decompositions and Applications”. In: *SIAM Review* 51.3 (Sept. 2009), pp. 455–500. DOI: 10.1137/07070111X.
- [KDS08] W. P. Krijnen, T. K. Dijkstra, and A. Stegeman. “On the non-existence of optimal solutions and the occurrence of “degeneracy” in the Candecomp/Parafac model”. In: *Psychometrika* 73.3 (2008), pp. 431–439.
- [Ken90] W. S. Kendall. “Probability, Convexity, and Harmonic Maps with Small Image I: Uniqueness and Fine Existence”. In: *Proc. London Math. Soc.* s3-61.2 (Sept. 1990), pp. 371–406. ISSN: 0024-6115. DOI: 10.1112/plms/s3-61.2.371.
- [KH20] T. G. Kolda and D. Hong. “Stochastic Gradients for Large-Scale Tensor Decomposition”. In: *SIAM Journal on Mathematics of Data Science* 2.4 (2020), pp. 1066–1095. DOI: 10.1137/19M1266265.
- [KHL89] J. B. Kruskal, R. A. Harshman, and M. E. Lundy. “How 3-MFA data can cause degenerate parafac solutions, among other relationships”. In: *Multisway Data Analysis*. NLD: North-Holland Publishing Co., 1989, pp. 115–122. ISBN: 0444874100.
- [KIS21] H. Knut, M. Irina, and L. F. Silva. “A Lagrangian approach to extremal curves on Stiefel manifolds”. In: *Journal of Geometric Mechanics* 13.1 (2021), pp. 55–72. ISSN: 1941-4889. DOI: 10.3934/jgm.2020031.
- [KL07] O. Koch and C. Lubich. “Dynamical Low-Rank Approximation”. In: *SIAM Journal on Matrix Analysis and Applications* 29.2 (2007), pp. 434–454. DOI: 10.1137/050639703.
- [KL10] O. Koch and C. Lubich. “Dynamical Tensor Approximation”. In: *SIAM Journal on Matrix Analysis and Applications* 31.5 (2010), pp. 2360–2375. DOI: 10.1137/09076578X.
- [KLW16] E. Kieri, C. Lubich, and H. Walach. “Discretized dynamical low-rank approximation in the presence of small singular values”. In: *SIAM Journal on Numerical Analysis* 54.2 (2016), pp. 1020–1038.
- [KSV14] D. Kressner, M. Steinlechner, and B. Vandereycken. “Low-rank tensor completion by Riemannian optimization”. In: *BIT Numerical Mathematics* 54.2 (2014), pp. 447–468.
- [Lan12] J. Landsberg. *Tensors: Geometry and Applications*. Graduate studies in mathematics. American Mathematical Society, 2012. ISBN: 9780821884812.

- [Lee13] J. M. Lee. *Introduction to Smooth Manifolds*. Graduate texts in mathematics: 218. Springer, 2013. ISBN: 9781441999818.
- [Lee18] J. M. Lee. *Introduction to Riemannian Manifolds*. Graduate texts in mathematics: 176. Springer, 2018. ISBN: 9783319917542.
- [LL17] D. Lazar and L. Lin. “Scale and curvature effects in principal geodesic analysis”. In: *Journal of Multivariate Analysis* 153 (2017), pp. 64–82. ISSN: 0047-259X. DOI: 10.1016/j.jmva.2016.09.009.
- [LO14] C. Lubich and I. V. Oseledets. “A projector-splitting integrator for dynamical low-rank approximation”. In: *BIT Numerical Mathematics* 54 (1 2014), pp. 171–188. DOI: 10.1007/s10543-013-0454-0.
- [Lub+13] C. Lubich, T. Rohwedder, R. Schneider, and B. Vandereycken. “Dynamical Approximation by Hierarchical Tucker and Tensor-Train Tensors”. In: *SIAM J. Matrix Anal. Appl.* 34.2 (Jan. 2013), pp. 470–494. ISSN: 1095-7162. DOI: 10.1137/120885723.
- [MA20] E. Massart and P.-A. Absil. “Quotient Geometry with Simple Geodesics for the Manifold of Fixed-Rank Positive-Semidefinite Matrices”. In: *SIAM Journal on Matrix Analysis and Applications* 41.1 (2020), pp. 171–198. DOI: 10.1137/18M1231389.
- [Mas80] J. Mason. “Near-best multivariate approximation by Fourier series, Chebyshev series and Chebyshev interpolation”. In: *Journal of Approximation Theory* 28.4 (1980), pp. 349–358. ISSN: 0021-9045. DOI: 10.1016/0021-9045(80)90069-6.
- [MB94] B. C. Mitchell and D. S. Burdick. “Slowly converging PARAFAC sequences: swamps and two-factor degeneracies”. In: *Journal of Chemometrics* 8.2 (1994), pp. 155–168.
- [MCB24] N. Mankovich, G. Camps-Valls, and T. Birdal. “Fun with Flags: Robust Principal Directions via Flag Manifolds”. In: *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, June 2024, pp. 330–340. DOI: 10.1109/CVPR52733.2024.00039.
- [MM24] A. Massarenti and M. Mella. “Bronowski’s conjecture and the identifiability of projective varieties”. In: *Duke Mathematical Journal* 173.17 (2024), pp. 3293–3316.
- [Moa02] M. Moakher. “Means and averaging in the group of rotations”. In: *SIAM Journal on Matrix Analysis and Applications* 24.1 (2002), pp. 1–16.

- [Moa06] M. Moakher. “On the Averaging of Symmetric Positive-Definite Tensors”. In: *Journal of Elasticity* 82 (3 Mar. 1, 2006), pp. 273–296. DOI: 10.1007/s10659-005-9035-z.
- [Moo16] C. Moosmüller. “ C^1 Analysis of Hermite Subdivision Schemes on Manifolds”. In: *SIAM Journal on Numerical Analysis* 54.5 (2016), pp. 3003–3031. DOI: 10.1137/15M1033459.
- [Moo17] C. Moosmüller. “Hermite subdivision on manifolds via parallel transport”. In: *Advances in Computational Mathematics* 43 (5 Oct. 1, 2017), pp. 1059–1074. DOI: 10.1007/s10444-017-9516-1.
- [MQZ14] H. Z. Munthe-Kaas, G. R. W. Quispel, and A. Zanna. “Symmetric spaces and Lie triple systems in numerical analysis of differential equations”. In: *BIT Numerical Mathematics* 54 (1 Mar. 1, 2014), pp. 257–282. DOI: 10.1007/s10543-014-0473-5.
- [Mun99] H. Munthe-Kaas. “High order Runge-Kutta methods on manifolds”. In: *Applied Numerical Mathematics* 29.1 (1999). Proceedings of the NSF/CBMS Regional Conference on Numerical Analysis of Hamiltonian Differential Equations, pp. 115–127. ISSN: 0168-9274. DOI: 10.1016/S0168-9274(98)00030-0.
- [MV03] C. Moler and C. Van Loan. “Nineteen Dubious Ways to Compute the Exponential of a Matrix, Twenty-Five Years Later”. In: *SIAM Review* 45.1 (2003), pp. 3–49. DOI: 10.1137/S00361445024180.
- [MV15] H. Munthe-Kaas and O. Verdier. “Integrators on Homogeneous Spaces: Isotropy Choice and Connections”. In: *Foundations of Computational Mathematics* 16.4 (July 2015), pp. 899–939. ISSN: 1615-3383. DOI: 10.1007/s10208-015-9267-7.
- [MW08] S. J. A. Malham and A. Wiese. “Stochastic Lie Group Integrators”. In: *SIAM Journal on Scientific Computing* 30.2 (2008), pp. 597–617. DOI: 10.1137/060666743.
- [MZ97] H. Munthe-Kaas and A. Zanna. “Numerical Integration of Differential Equations on Homogeneous Manifolds”. In: *Foundations of Computational Mathematics*. Ed. by F. Cucker and M. Shub. Berlin, Heidelberg: Springer Berlin Heidelberg, 1997, pp. 305–315. ISBN: 978-3-642-60539-0.
- [MZM25] S. Mataigne, R. Zimmermann, and N. Miolane. “An Efficient Algorithm for the Riemannian Logarithm on the Stiefel Manifold for a Family of Riemannian Metrics”. In: *SIAM Journal on Matrix Analysis and Applications* 46.2 (2025), pp. 879–905. DOI: 10.1137/24M1647801.

- [NB13] F. Nielsen and R. Bhatia. *Matrix Information Geometry*. Berlin, Heidelberg: Springer, 2013. ISBN: 978-3-642-30231-2. DOI: 10.1007/978-3-642-30232-9.
- [NDK08] C. Navasca, L. De Lathauwer, and S. Kindermann. “Swamp reducing technique for tensor decomposition”. In: *2008 16th european signal processing conference*. IEEE. 2008, pp. 1–5.
- [NST18] Y. Nakatsukasa, O. Sète, and L. N. Trefethen. “The AAA Algorithm for Rational Approximation”. In: *SIAM Journal on Scientific Computing* 40.3 (Jan. 2018), A1494–A1522. ISSN: 1095-7197. DOI: 10.1137/16m1106122.
- [ONe66] B. O’Neill. “The fundamental equations of a submersion.” In: *Michigan Mathematical Journal* 13.4 (1966), pp. 459–469. DOI: 10.1307/mmj/1028999604.
- [ONe83] B. O’Neill. *Semi-Riemannian Geometry With Applications to Relativity*. en. Academic Press, July 1983. ISBN: 978-0-08-057057-0.
- [ORU18] I. V. Oseledets, M. V. Rakhuba, and A. Uschmajew. “Alternating Least Squares as Moving Subspace Correction”. In: *SIAM Journal on Numerical Analysis* 56.6 (2018), pp. 3459–3479. DOI: 10.1137/17M1148712.
- [ORU23] I. V. Oseledets, M. V. Rakhuba, and A. Uschmajew. “Local convergence of alternating low-rank optimization methods with overrelaxation”. In: *Numerical Linear Algebra with Applications* 30.3 (2023), e2459.
- [Ose11] I. V. Oseledets. “Tensor-Train Decomposition”. In: *SIAM Journal on Scientific Computing* 33.5 (2011), pp. 2295–2317. DOI: 10.1137/090752286.
- [Ose13] I. V. Oseledets. “Constructive representation of functions in low-rank tensor formats”. In: *Constructive Approximation* 37.1 (2013), pp. 1–18.
- [Osi25] A. Osinsky. “Close to optimal column approximation using a single SVD”. In: *Linear Algebra and its Applications* 725 (2025), pp. 359–377. ISSN: 0024-3795. DOI: 10.1016/j.laa.2025.07.016.
- [OT10] I. Oseledets and E. Tyrtyshnikov. “TT-cross approximation for multidimensional arrays”. In: *Linear Algebra and its Applications* 432.1 (2010), pp. 70–88. ISSN: 0024-3795. DOI: 10.1016/j.laa.2009.07.024.
- [Paa00] P. Paatero. “Construction and analysis of degenerate PARAFAC models”. In: *Journal of Chemometrics: A Journal of the Chemometrics Society* 14.3 (2000), pp. 285–299.

- [Pen20] X. Pennec. “3 - Manifold-valued image processing with SPD matrices”. In: *Riemannian Geometric Statistics in Medical Image Analysis*. Ed. by X. Pennec, S. Sommer, and T. Fletcher. Cambridge, MA: Academic Press, 2020, pp. 75–134. ISBN: 978-0-12-814725-2. DOI: 10.1016/B978-0-12-814725-2.00010-8.
- [Pet16] P. Petersen. *Riemannian Geometry*. New York, NY: Springer Cham, 2016. ISBN: 978-3-319-26652-7. DOI: 10.1007/978-3-319-26654-1.
- [PFA06] X. Pennec, P. Fillard, and N. Ayache. “A Riemannian Framework for Tensor Computing”. In: *International Journal of Computer Vision* 66 (1 Jan. 1, 2006), pp. 41–66. DOI: 10.1007/s11263-005-3222-z.
- [PM19] A. Petersen and H.-G. Müller. “Fréchet regression for random objects with Euclidean predictors”. In: *The Annals of Statistics* 47.2 (2019), pp. 691–719. DOI: 10.1214/17-AOS1624.
- [PTC13] A.-H. Phan, P. Tichavsky, and A. Cichocki. “Low complexity damped Gauss–Newton algorithms for CANDECOMP/PARAFAC”. In: *SIAM Journal on Matrix Analysis and Applications* 34.1 (2013), pp. 126–147.
- [Ric66] J. R. Rice. “A Theory of Condition”. In: *SIAM Journal on Numerical Analysis* 3.2 (1966), pp. 287–310. DOI: 10.1137/0703023. eprint: 10.1137/0703023.
- [RS05] M. Reid and B. Szendroi. *Geometry and Topology*. Cambridge University Press, 2005. DOI: 10.1017/CB09780511807510.
- [RTW22] D. Rubin, A. Townsend, and H. Wilber. “Bounding Zolotarev numbers using Faber rational functions”. In: *Constructive Approximation* 56.2 (2022), pp. 207–232.
- [RU13] T. Rohwedder and A. Uschmajew. “On Local Convergence of Alternating Schemes for Optimization of Convex Problems in the Tensor Train Format”. In: *SIAM Journal on Numerical Analysis* 51.2 (2013), pp. 1134–1162. DOI: 10.1137/110857520.
- [Rya02] R. A. Ryan. *Introduction to tensor products of Banach spaces*. Vol. 73. Springer, 2002.
- [Sai16] A. K. Saibaba. “HOID: Higher order interpolatory decomposition for tensors based on Tucker representation”. In: *SIAM Journal on Matrix Analysis and Applications* 37.3 (2016), pp. 1223–1249.
- [Sam+12] C. Samir, P.-A. Absil, A. Srivastava, and E. Klassen. “A gradient-descent method for curve fitting on Riemannian manifolds”. In: *Foundations of Computational Mathematics* 12.1 (2012), pp. 49–73.

- [Sav14] D. V. Savostyanov. “Quasioptimality of maximum-volume cross interpolation of tensors”. In: *Linear Algebra and its Applications* 458 (2014), pp. 217–244.
- [Sch69] M. H. Schultz. “ L^∞ -multivariate approximation theory”. In: *SIAM Journal on Numerical Analysis* 6.2 (1969), pp. 161–183.
- [SCK24] A. Séguin, G. Ceruti, and D. Kressner. “From low-rank retractions to dynamical low-rank approximation and back”. In: *BIT Numerical Mathematics* 64 (3 2024). DOI: 10.1007/s10543-024-01028-7.
- [SCW23] N. Sharon, R. S. Cohen, and H. Wendland. “On Multiscale Quasi-Interpolation of Scattered Scalar- and Manifold-Valued Functions”. In: *SIAM Journal on Scientific Computing* 45.5 (2023), A2458–A2482. DOI: 10.1137/22M1528306.
- [Shi18] Y. Shitov. “A Counterexample to Comon’s Conjecture”. In: *SIAM Journal on Applied Algebra and Geometry* 2.3 (2018), pp. 428–443. DOI: 10.1137/17M1131970.
- [Sid+17] N. D. Sidiropoulos, L. De Lathauwer, X. Fu, K. Huang, E. E. Papalexakis, and C. Faloutsos. “Tensor Decomposition for Signal Processing and Machine Learning”. In: *IEEE Transactions on Signal Processing* 65.13 (2017), pp. 3551–3582. DOI: 10.1109/TSP.2017.2690524.
- [SK23] C. Strössner and D. Kressner. “Low-Rank Tensor Approximations for Solving Multimarginal Optimal Transport Problems”. In: *SIAM Journal on Imaging Sciences* 16.1 (2023), pp. 169–191. DOI: 10.1137/22M1478355.
- [SL08] V. de Silva and L.-H. Lim. “Tensor Rank and the Ill-Posedness of the Best Low-Rank Approximation Problem”. In: *SIAM Journal on Matrix Analysis and Applications* 30.3 (2008), pp. 1084–1127. DOI: 10.1137/06066518X.
- [SL10] B. Savas and L.-H. Lim. “Quasi-Newton methods on Grassmannians and multilinear approximations of tensors”. In: *SIAM Journal on Scientific Computing* 32.6 (2010), pp. 3352–3393.
- [Smi+04] K. Smith, L. Kahanpää, P. Kekäläinen, and W. Traves. *An invitation to algebraic geometry*. Springer Science & Business Media, 2004.
- [SOF23] A. Saiapin, I. Oseledets, and E. Frolov. “Dynamic Collaborative Filtering for Matrix-and Tensor-based Recommender Systems”. In: *arXiv preprint arXiv:2312.10064* (2023).

- [SSK20] K. M. Soodhalter, E. de Sturler, and M. E. Kilmer. “A survey of subspace recycling iterative methods”. In: *GAMM-Mitteilungen* 43.4 (2020), e202000016. DOI: <https://doi.org/10.1002/gamm.202000016>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/gamm.202000016>.
- [SSK24] C. Strössner, B. Sun, and D. Kressner. “Approximation in the extended functional tensor train format”. In: *Advances in Computational Mathematics* 50.3 (2024), p. 54.
- [Ste06] A. Stegeman. “Degeneracy in Candecomp/Parafac explained for $p \times p \times 2$ arrays of rank $p + 1$ or higher”. In: *Psychometrika* 71.3 (2006), pp. 483–501.
- [Ste07] A. Stegeman. “Degeneracy in Candecomp/Parafac and Indscal explained for several three-sliced arrays with a two-valued typical rank”. In: *Psychometrika* 72.4 (2007), pp. 601–619.
- [Ste08] A. Stegeman. “Low-Rank Approximation of Generic $p \times q \times 2$ Arrays and Diverging Components in the Candecomp/Parafac Model”. In: *SIAM Journal on Matrix Analysis and Applications* 30.3 (2008), pp. 988–1007. DOI: 10.1137/050644677.
- [SVD13] L. Sorber, M. Van Barel, and L. De Lathauwer. “Optimization-Based Algorithms for Tensor Decompositions: Canonical Polyadic Decomposition, Decomposition in Rank-(Lr,Lr,1) Terms, and a New Generalization”. In: *SIAM Journal on Optimization* 23.2 (2013), pp. 695–720. DOI: 10.1137/120868323.
- [SvV22] L. Swijsen, J. van der Veken, and N. Vannieuwenhoven. “Tensor completion using geodesics on Segre manifolds”. In: *Numerical Linear Algebra with Applications* 29 (Apr. 2022). DOI: 10.1002/nla.2446.
- [Swi22] L. Swijsen. “Tensor decompositions and Riemannian optimization”. PhD thesis. KU Leuven, 2022.
- [TB06] G. Tomasi and R. Bro. “A comparison of algorithms for fitting the PARAFAC model”. In: *Computational Statistics & Data Analysis* 50.7 (2006), pp. 1700–1734. ISSN: 0167-9473. DOI: 10.1016/j.csda.2004.11.013.
- [TC24] A. Taveira Blomenhofer and A. Casarotti. “Nondefectivity of invariant secant varieties”. In: *arXiv:2312.12335v2* (2024).
- [Tem18] V. Temlyakov. *Multivariate Approximation*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge, England: Cambridge University Press, 2018. DOI: 10.1017/9781108689687.

- [Tho13] P. Thomas Fletcher. “Geodesic Regression and the Theory of Least Squares on Riemannian Manifolds”. In: *International Journal of Computer Vision* 105 (2 2013), pp. 171–185. DOI: 10.1007/s11263-012-0591-y.
- [TP21] Y. Thanwerdas and X. Pennec. “Geodesics and Curvature of the Quotient-Affine Metrics on Full-Rank Correlation Matrices”. In: *Geometric Science of Information*. Ed. by F. Nielsen and F. Barbaresco. Cham: Springer International Publishing, 2021, pp. 93–102. ISBN: 978-3-030-80209-7.
- [TP23] Y. Thanwerdas and X. Pennec. “O(n)-invariant Riemannian metrics on SPD matrices”. In: *Linear Algebra and its Applications* 661 (2023), pp. 163–201. ISSN: 0024-3795. DOI: 10.1016/j.laa.2022.12.009.
- [Tre19] L. N. Trefethen. *Approximation Theory and Approximation Practice, Extended Edition*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2019. DOI: 10.1137/1.9781611975949.
- [TT13] A. Townsend and L. N. Trefethen. “An extension of Chebfun to two dimensions”. In: *SIAM Journal on Scientific Computing* 35.6 (2013), pp. C495–C518.
- [Usc12] A. Uschmajew. “Local convergence of the alternating least squares algorithm for canonical tensor approximation”. In: *SIAM Journal on Matrix Analysis and Applications* 33.2 (2012), pp. 639–652.
- [UT19] M. Udell and A. Townsend. “Why are big data matrices approximately low rank?” In: *SIAM Journal on Mathematics of Data Science* 1.1 (2019), pp. 144–160.
- [UV13] A. Uschmajew and B. Vandereycken. “The geometry of algorithms using hierarchical tensors”. In: *Linear Algebra and its Applications* 439.1 (2013), pp. 133–166. ISSN: 0024-3795. DOI: 10.1016/j.laa.2013.03.016.
- [UV20] A. Uschmajew and B. Vandereycken. “Geometric Methods on Low-Rank Matrix and Tensor Manifolds”. In: *Handbook of Variational Methods for Nonlinear Geometric Data*. Ed. by P. Grohs, M. Holler, and A. Weinmann. Cham: Springer International Publishing, 2020, pp. 261–313. ISBN: 978-3-030-31351-7. DOI: 10.1007/978-3-030-31351-7_9.
- [VAV12] B. Vandereycken, P.-A. Absil, and S. Vandewalle. “A Riemannian geometry with complete geodesics for the set of positive semidefinite matrices of fixed rank”. In: *IMA Journal of Numerical Analysis* 33.2 (July 2012), pp. 481–514. ISSN: 0272-4979. DOI: 10.1093/imanum/drs006.

- [Ven18] F. Venturelli. *Prehomogeneous tensor spaces*. 2018. arXiv: 1606.07257 [math.AG].
- [VVD25] C. Vermeylen, N. Vervliet, and L. De Lathauwer. “Reducing swamp behavior for the canonical polyadic decomposition problem by rank-1 freezing”. In: *Numerical Algorithms* (2025), pp. 1–29.
- [VVM12] N. Vannieuwenhoven, R. Vandebril, and K. Meerbergen. “A New Truncation Strategy for the Higher-Order Singular Value Decomposition”. In: *SIAM Journal on Scientific Computing* 34.2 (2012), A1027–A1052. DOI: 10.1137/110836067.
- [Wan+25] H. Wang, R. Vandebril, J. V. der Veken, and N. Vannieuwenhoven. *Manifold-valued function approximation from multiple tangent spaces*. 2025. arXiv: 2504.12892 [math.NA].
- [WC14] L. Wang and M. T. Chu. “On the global convergence of the alternating least squares method for rank-one approximation to generic tensors”. In: *SIAM Journal on Matrix Analysis and Applications* 35.3 (2014), pp. 1058–1072.
- [WC88] G. Wu and W. Chen. “A matrix inequality and its geometric applications”. In: *Acta Math. Sinica* 31.3 (1988), pp. 348–355.
- [WCY15] L. Wang, M. T. Chu, and B. Yu. “Orthogonal low rank tensor approximation: alternating least squares method and its global convergence”. In: *SIAM Journal on Matrix Analysis and Applications* 36.1 (2015), pp. 1–19.
- [Wen04] H. Wendland. “Moving least squares”. In: *Scattered Data Approximation*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge, England: Cambridge University Press, 2004, pp. 35–45.
- [Yan95] B. Yang. “Projection approximation subspace tracking”. In: *IEEE Transactions on Signal Processing* 43.1 (1995), pp. 95–107. DOI: 10.1109/78.365290.
- [YWL22] K. Ye, K. S.-W. Wong, and L.-H. Lim. “Optimization on flag manifolds”. In: *Mathematical Programming* 194.1 (2022), pp. 621–660.
- [ZB24] R. Zimmermann and R. Bergmann. “Multivariate Hermite Interpolation on Riemannian Manifolds”. In: *SIAM Journal on Scientific Computing* 46.2 (2024), A1276–A1297. DOI: 10.1137/22M1541071.
- [ZH22] R. Zimmermann and K. Hüper. “Computing the Riemannian Logarithm on the Stiefel Manifold: Metrics, Methods, and Performance”. In: *SIAM Journal on Matrix Analysis and Applications* 43.2 (2022), pp. 953–980. DOI: 10.1137/21M1425426.

- [Zim17] R. Zimmermann. “A Matrix-Algebraic Algorithm for the Riemannian Logarithm on the Stiefel Manifold under the Canonical Metric”. In: *SIAM Journal on Matrix Analysis and Applications* 38.2 (2017), pp. 322–342. DOI: 10.1137/16M1074485.
- [Zim20] R. Zimmermann. “Hermite Interpolation and Data Processing Errors on Riemannian Matrix Manifolds”. In: *SIAM Journal on Scientific Computing* 42.5 (Jan. 2020), A2593–A2619. ISSN: 1095-7197. DOI: 10.1137/19m1282878.
- [ZN19] E. Zhang and L. Noakes. “The cubic de Casteljau construction and Riemannian cubics”. In: *Computer Aided Geometric Design* 75 (2019), p. 101789. ISSN: 0167-8396. DOI: 10.1016/j.cagd.2019.101789.
- [ZS25] R. Zimmermann and J. Stoye. “High Curvature Means Low Rank: On the Sectional Curvature of Grassmann and Stiefel Manifolds and the Underlying Matrix Trace Inequalities”. In: *SIAM Journal on Matrix Analysis and Applications* 46.1 (2025), pp. 748–779. DOI: 10.1137/24M1655755.

Statement on the use of Generative AI

I did not use generative AI assistance tools during the research/writing process of my thesis, except for mere language assistance.

The text, code, and images in this thesis are my own (unless otherwise specified). Generative AI has only been used in accordance with the KU Leuven guidelines and appropriate references have been added. I have reviewed and edited the content as needed and I take full responsibility for the content of the thesis.

FACULTY OF ENGINEERING SCIENCE
DEPARTMENT OF COMPUTER SCIENCE
Celestijnenlaan 200A box 2402
B-3001 Leuven

