Modeling and Sampling of Spectrally Structured Signals

Elvander, Filip

2020

Document Version:
Publisher's PDF, also known as Version of record

– CENTRUM SCIENTIARUM MATHEMATICARUM –

Modeling and Sampling of Spectrally Structured Signals

FILIP ELVANDER

Lund University
Faculty of Engineering
Centre for Mathematical Sciences
Mathematical Statistics

Printed by Media-Tryck, Lund 2020 NORDIC SWAN ECOLABEL 3041 0903

LUND UNIVERSITY

LTH
FACULTY OF ENGINEERING

Doctoral Theses in Mathematical Sciences 2020:2

ISBN 978-91-7895-488-9
ISRN LUTFMS-1048-2020
ISSN 1404-0034

Modeling and Sampling of Spectrally Structured Signals

av

Filip Elvander

Lunds universitet

LTH

Lunds tekniska högskola

Avhandling som för avläggande av teknologie doktorsexamen vid tekniska fakulteten vid Lunds universitet kommer att offentligen försvaras i sal MH:Riesz, Matematikcentrum, Sölvegatan 18, fredagen den 12 juni 2020 kl. 9^{00}.

Fakultetsopponent är universitetslektor Esa Ollila, Aalto-universitetet, Esbo, Finland.

Organization
LUND UNIVERSITY

Centre for Mathematical Sciences
Mathematical Statistics
Box 118
SE-221 00 Lund, Sweden

Author
Filip Elvander

Document name
DOCTORAL DISSERTATION

Date of disputation
2020-06-12

Sponsoring organization

Title
Modeling and Sampling of Spectrally Structured Signals

Abstract
This thesis consists of five papers concerned with the modeling of stochastic signals, as well as deterministic signals in stochastic noise, exhibiting different kinds of structure. This structure is manifested as the existence of finite-dimensional parameterizations, and/or in the geometry of the signals' spectral representations.

The two first papers of the thesis, Papers A and B, consider the modeling of differences, or distances, between stochastic processes based on their second-order statistics, i.e., covariances. By relating the covariance structure of a stochastic process to spectral representations, it is proposed to quantify the dissimilarity between two processes in terms of the cost associated with morphing one spectral representation to the other, with the cost of morphing being defined in terms of the solutions to optimal mass transport problems. The proposed framework allows for modeling smooth changes in the frequency characteristics of stochastic processes, which is shown to yield interpretable and physically sensible predictions when used in applications of temporal and spatial spectral estimation. Also presented are efficient computational tools, allowing for the framework to be used in high-dimensional problems.

Paper C considers the modeling of so-called inharmonic signals, i.e., signals that are almost, but not quite, harmonic. Such signals appear in many fields of signal processing, not least in audio. Inharmonicity may be interpreted as a deviation from a parametric structure, as well as from a particular spectral structure. Based on these views, as well as on a third, stochastic interpretation, Paper C proposes three different definitions of the concept of fundamental frequency for inharmonic signals, and studies the estimation theoretical implications of utilizing either of these definitions. Paper D then considers deliberate deviations from a parametric signal structure arising in spectroscopy applications. With the motivation of decreasing the computational complexity of parameter estimation, the paper studies the implications of estimating the parameters of the signal in a sequential fashion, starting out with a simplified model that is then refined step by step.

Lastly, Paper E studies how parametric descriptions of signals can be leveraged as to design optimal, in an estimation theoretical sense, schemes for sampling or collecting measurements from the signal. By means of a convex program, samples are selected as to minimize bounds on estimator variance, allowing for efficiently measuring parametric signals, even when the parametrization is only partially known.

Key words
Spectral estimation, Parameter estimation, Optimal mass transport, Covariance interpolation, Misspecified models, Inharmonicity

Classification system and/or index terms (if any)

Supplementary bibliographical information

Language
English

ISSN and key title
1404-0034, Doctoral Theses in Mathematical Sciences 2020:2

ISBN
978-91-7895-488-9

Recipient's notes

Number of pages
256

ISRN
LUTFMS-1048-2020

Security classification

Modeling and Sampling of Spectrally Structured Signals

Filip Elvander

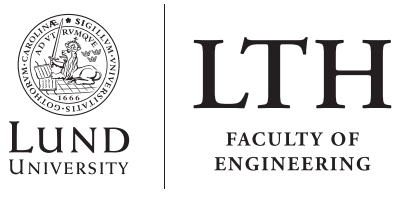

Faculty of Engineering
Centre for Mathematical Sciences
Mathematical Statistics

Mathematical Statistics
Centre for Mathematical Sciences
Lund University
Box 118
SE-221 00 Lund
Sweden

http://www.maths.lu.se/

Doctoral Theses in Mathematical Sciences 2020:2
ISSN 1404-0034

ISBN 978-91-7895-488-9 (print)
ISBN 978-91-7895-489-6 (pdf)
LUTFMS-1048-2020

Printed in Sweden by Media-Tryck, Lund 2020

Contents

Abstract

This thesis consists of five papers concerned with the modeling of stochastic signals, as well as deterministic signals in stochastic noise, exhibiting different kinds of structure. This structure is manifested as the existence of finite-dimensional parameterizations, and/or in the geometry of the signals' spectral representations. The two first papers of the thesis, Papers A and B, consider the modeling of differences, or distances, between stochastic processes based on their second-order statistics, i.e., covariances. By relating the covariance structure of a stochastic process to spectral representations, it is proposed to quantify the dissimilarity between two processes in terms of the cost associated with morphing one spectral representation to the other, with the cost of morphing being defined in terms of the solutions to optimal mass transport problems. The proposed framework allows for modeling smooth changes in the frequency characteristics of stochastic processes, which is shown to yield interpretable and physically sensible predictions when used in applications of temporal and spatial spectral estimation. Also presented are efficient computational tools, allowing for the framework to be used in high-dimensional problems. Paper C considers the modeling of so-called inharmonic signals, i.e., signals that are almost, but not quite, harmonic. Such signals appear in many fields of signal processing, not least in audio. Inharmonicity may be interpreted as a deviation from a parametric structure, as well as from a particular spectral structure. Based on these views, as well as on a third, stochastic interpretation, Paper C proposes three different definitions of the concept of fundamental frequency for inharmonic signals, and studies the estimation theoretical implications of utilizing either of these definitions. Paper D then considers deliberate deviations from a parametric signal structure arising in spectroscopy applications. With the motivation of decreasing the computational complexity of parameter estimation, the paper studies the implications of estimating the parameters of the signal in a sequential fashion, starting out with a simplified model that is then refined step by step. Lastly, Paper E studies how parametric descriptions of signals can be leveraged as to design optimal, in an estimation theoretical sense, schemes for sampling or collecting measurements from the signal. By means of a convex program, samples are selected as to minimize bounds on estimator variance, allow-

ing for efficiently measuring parametric signals, even when the parametrization is only partially known.

Populärvetenskaplig sammanfattning på svenska

Den här avhandlingen undersöker sätt att modellera, eller beskriva, signaler som uppvisar vissa former av struktur. Många signaler som vi möter i vårt vardagliga liv uppvisar mönster som gör det möjligt för oss att utvinna information ur dem. Exempelvis tolkar vi sekvensen av ljud i mänskligt tal som ord, och klang och tonhöjd låter oss identifiera vem det är som talar. Detta möjliggörs av att det finns mönster i ljudets frekvenser, eller specifikt i hur röstens energi är fördelad över olika frekvensband. En sådan energifördelning kallas för spektrum. Liknande strukturer dyker upp i signaler från så skilda områden som radar, sonar och mobilkommunikation. I avhandlingen presenteras verktyg som utnyttjar sådan struktur för att mäta avstånd mellan signaler och därmed kunna identifiera och särskilja dem. I de två första artiklarna, och delvis i den tredje, kvantifieras avståndet mellan två signaler av hur kostsamt det är att fördela om deras energi för att deras spektra ska bli identiska. Om signalerna är mänskliga röster innebär detta att de efter omfördelningen av energi kommer att låta likadant. Att mäta avstånd på det här sättet får till följd att signaler bedöms som mer olika ju längre ifrån varandra i frekvens deras spektra ligger. Intuitivt är detta mycket rimligt: två barnaröster är mer lika varandra än rösten hos en vuxen man, även om de inte är helt identiska. Avståndsmåttet möjliggör även skapandet av syntetiska, mellanliggande spektra; det går alltså att skapa ”medelröster” som ligger mitt emellan två personer. När samma typ av verktyg används i radarapplikationer blir följden att man kan följa och förutse rörelser hos exempelvis bilar och flygplan.

Avhandlingen behandlar också följderna av att ignorera en viss del av strukturen hos en signal. Anledningen till att göra detta är att det i vissa sammanhang är beräkningsmässigt kostsamt att ta hänsyn till en detaljerad beskrivning av en signal. I den fjärde artikeln undersöks hur väl man kan uppskatta en signals egenskaper om man endast tar hänsyn till en del av dess struktur.

I avhandlingens femte och sista artikel undersöks hur en signals struktur kan utnyttjas för att avgöra hur den på bästa sätt bör mätas. I vissa sammanhang är det dyrt eller tidskrävande av samla in mätningar av en signal; en enda mätpunkt kan

kräva att ett helt experiment utförs. Detta får till följd att det i områden såsom spektroskopi är praktiskt omöjligt att samla in mätningar uniformt, alltså för alla kombinationer av experimentinställningar. För att angripa detta problem föreslår avhandlingens femte artikel en metod för att välja ut ett litet antal mätningar på ett sådant sätt att signalens egenskaper kan skattas med en precision som är försumbart lägre än om alla mätningar kunnat samlas in. Metoden har visats leda till en tiofaldig reduktion av experimenttiden för visa typer av spektroskopiexperiment.

Acknowledgements

Firstly, I want to thank my supervisor Prof. Andreas Jakobsson; with your never-ending enthusiasm, supportiveness in all situations, and generosity in letting me explore topics of my own choice, I could not have asked for a better guide during these five years. I would also like to thank Assoc. Prof. Johan Karlsson for introducing me to the wonders of optimal mass transport and for shaping my research interests during the latter half of my PhD. A big thanks to Assoc. Prof. Toon van Waterschoot for hosting me in his group at KU Leuven during the fall of 2018 and for making me feel very much at home. To Johan Swärd, my long-term co-author and go-to discussion partner, thank you for making this time so much fun! Also, thanks to Isabel Haasler, my more recent co-author, for a rewarding collaboration as well as enjoyable conference trips. I would also like to thank the co-authors of papers not included in the thesis: Stefan Adalbjörnsson, Ted Kronvall, Maria Juhlin, Martin Jälmby, Randall Ali, Andreas Jansson, Göran Carlström, Mikael Akke, Jie Ding, and Shiwen Lei. I am also most grateful to the technical and administrative staff at Mathematical Statistics; thank you for resolving all strange and practical things.

Till min familj, Johan, Christine, Linnéa, Anna-Lovisa och Alexander, tack för att ni gör allting möjligt!

Lund, 2020 *Filip Elvander*

List of papers

This thesis is based on the following papers:

A Filip Elvander, Andreas Jakobsson, and Johan Karlsson, "Interpolation and Extrapolation of Toeplitz Matrices via Optimal Mass Transport". *IEEE Transactions on Signal Processing*, vol. 66, pp. 5285-5298, October 15, 2018.

B Filip Elvander, Isabel Haasler, Andreas Jakobsson, and Johan Karlsson, "Multi-marginal optimal transport using partial information with applications in robust localization and sensor fusion". *Signal Processing*, vol. 171, pp. 1-19, June 2020.

C Filip Elvander and Andreas Jakobsson, "Defining Fundamental Frequency for Almost Harmonic Signals". *Submitted*.

D Filip Elvander, Johan Swärd, and Andreas Jakobsson, "Mismatched Estimation of Polynomially Damped Signals". *2019 IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, pp. 246 - 250, Guadeloupe, France, December 2019.

E Johan Swärd, Filip Elvander, and Andreas Jakobsson, "Designing sampling schemes for multi-dimensional data". *Signal Processing*, vol. 150, pp. 1-10, September 2018.

Additional papers not included in the thesis:

1. Göran Carlström, Filip Elvander, Johan Swärd, Andreas Jakobsson, and Mikael Akke, "Rapid NMR Relaxation Measurements Using Optimal Non-Uniform Sampling of Multi-Dimensional Accordion Data Analyzed by a Sparse Reconstruction Method", *The Journal of Physical Chemistry A*, vol. 123, pp. 5718-5723, June 2019.

2. Filip Elvander, Johan Swärd, and Andreas Jakobsson, "Multi-dimensional grid-less estimation of saturated signals". *Signal Processing*, vol. 145, April 2018, pages 37-47.

3. Filip Elvander, Johan Swärd, and Andreas Jakobsson, "Online Estimation of Multiple Harmonic Signals". *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 2, pp. 273-284, February 2017.

4. Filip Elvander, Stefan Ingi Adalbjörnsson, Ted Kronvall, and Andreas Jakobsson, "An Adaptive Penalty Multi-Pitch Estimator with Self-Regularization". *Signal Processing*, vol. 127, pp. 56-70, October 2016.

5. Andreas Jansson, Filip Elvander, Peter Almers, and Andreas Jakobsson, "Range-based Radar Model Structure Selection", *27th European Signal Processing Conference*, 2020. Submitted.

6. Filip Elvander, Jie Ding, and Andreas Jakobsson, "On Harmonic Approximations of Inharmonic Signals", *45th IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2020. Accepted for publication in proceedings.

7. Filip Elvander, Randall Ali, Andreas Jakobsson, and Toon van Waterschoot, "Offline Noise Reduction Using Optimal Mass Transport Induced Covariance Interpolation", *27th European Signal Processing Conference*. A Coruña, Spain, September 2-6, 2019.

8. Maria Juhlin, Filip Elvander, and Andreas Jakobsson, and Toon van Waterschoot, "Defining Graph Signal Distances Using an Optimal Mass Transport Framework", *27th European Signal Processing Conference*. A Coruña, Spain, September 2-6, 2019.

9. Filip Elvander, Isabel Haasler, Andreas Jakobsson, and Johan Karlsson, "Non-Coherent Sensor Fusion via Entropy Regularized Optimal Mass Transport", *44th IEEE International Conference on Acoustics, Speech, and Signal Processing*, pp. 4415-4419. Brighton, UK, May 12-17, 2019.

10. Martin Jälmby, Johan Swärd, Filip Elvander, and Andreas Jakobsson, "Computationally Efficient Estimation of Multi-dimensional Damped Modes Using Sparse Wideband Dictionaries", *26th European Signal Processing Conference*, pp. 1759-1763. Rome, Italy, September 3-7, 2018.

11. Filip Elvander, Isabel Haasler, Andreas Jakobsson, and Johan Karlsson, "Tracking and Sensor Fusion in Direction of Arrival Estimation Using

Optimal Mass Transport", *26th European Signal Processing Conference*, pp. 1631-1635. Rome, Italy, September 3-7, 2018.

12. Maria Juhlin, Filip Elvander, Johan Swärd, and Andreas Jakobsson, "Fast Grid-less Estimation of Damped Modes", *2018 International Symposium on Intelligent Signal Processing and Communication Systems*, pp. 346-351. Ishigaki Island, Japan, November 27-30, 2018.

13. Filip Elvander, Andreas Jakobsson, and Johan Karlsson, "Using Optimal Mass Transport for Tracking and Interpolation of Toeplitz Covariance Matrices", *43rd IEEE International Conference on Acoustics, Speech, and Signal Processing*, pp. 4469 - 4473. Calgary, Canada, April 15-20, 2018.

14. Filip Elvander, Johan Swärd, and Andreas Jakobsson, "Grid-less Estimation of Saturated Signals", *51st Asilomar Conf. on Signals, Systems, and Computers*, pp. 372 - 376. Pacific Grove, USA, October 29 - November 1, 2017.

15. Johan Swärd, Filip Elvander, and Andreas Jakobsson, "Designing Optimal Sampling Schemes for Multi-Dimensional Data", *51st Asilomar Conf. on Signals, Systems, and Computers*, pp. 850 - 852. Pacific Grove, USA, October 29 - November 1, 2017.

16. Filip Elvander, Stefan Ingi Adalbjörnsson, Johan Karlsson, and Andreas Jakobsson, "Using Optimal Transport for Estimating Inharmonic Pitch Signals". *42nd IEEE International Conference on Acoustics, Speech, and Signal Processing*, pp. 331-335. New Orleans, USA, March 5-9, 2017.

17. Johan Swärd, Filip Elvander, and Andreas Jakobsson, "Designing Optimal Sampling Schemes", *25th European Signal Processing Conference*, pp. 912 - 916. Kos island, Greece, August 28 - September 2, 2017.

18. Filip Elvander, Johan Swärd, and Andreas Jakobsson, "Time-Recursive Multi-Pitch Estimation Using Group Sparse Recursive Least Squares", *50th Annual Asilomar Conference on Signals, Systems, and Computers*, pp. 369-373. Pacific Grove, USA, November 6-9, 2016.

19. Shiwen Lei, Filip Elvander, Johan Swärd, Stefan Ingi Adalbjörnsson, and Andreas Jakobsson, "Computationally Efficient Multi-Pitch Estimation Using Sparsity", *11th IMA International Conference on Mathematics in Signal Processing*, pp. 1-4. Birmingham, UK, November 12-14, 2016.

20. Filip Elvander, Stefan Ingi Adalbjörnsson, and Andreas Jakobsson, "Robust Non-Negative Least Squares Using Sparsity", *24th European Signal Processing Conference*, pp. 61-65. Budapest, Hungary, August 29 - September 2, 2016.

21. Ted Kronvall, Filip Elvander, Stefan Ingi Adalbjörnsson, and Andreas Jakobsson, "Multi-Pitch Estimation via Fast Group Sparse Learning", *24th European Signal Processing Conference*, pp. 1093 - 1097. Budapest, Hungary, August 29 - September 2, 2016.

22. Ted Kronvall, Filip Elvander, Stefan Ingi Adalbjörnsson, and Andreas Jakobsson, "An Adaptive Penalty Approach to Multi-Pitch Estimation", *23rd European Signal Processing Conference*, pp. 31-35. Nice, France, August 31 - September 4, 2015.

Author's contributions

Filip Elvander's contributions to each of the papers included in the Ph.D thesis are mentioned in bold italics below each paper title.

Paper A: Filip Elvander, Andreas Jakobsson, and Johan Karlsson, "Interpolation and Extrapolation of Toeplitz Matrices via Optimal Mass Transport". *IEEE Transactions on Signal Processing*, vol. 66, pp. 5285-5298, October 15, 2018.

Joint main idea, derivations, numerical simulations, and write-up.

Paper B: Filip Elvander, Isabel Haasler, Andreas Jakobsson, and Johan Karlsson, "Multi-marginal optimal transport using partial information with applications in robust localization and sensor fusion". *Signal Processing*, vol. 171, pp. 1-19, June 2020.

Joint main idea, derivations, numerical simulations, and write-up.

Paper C: Filip Elvander and Andreas Jakobsson, "Defining Fundamental Frequency for Almost Harmonic Signals". *Submitted*.

Joint main idea, derivations, numerical simulations, and write-up.

Paper D: Filip Elvander, Johan Swärd, and Andreas Jakobsson, "Mismatched Estimation of Polynomially Damped Signals". *2019 IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, pp. 246 - 250, Guadeloupe, France, December 2019.

Joint main idea, derivations, numerical simulations, and write-up.

Paper E: Johan Swärd, Filip Elvander, and Andreas Jakobsson, "Designing sampling schemes for multi-dimensional data". *Signal Processing*, vol. 150, pp. 1-10, September 2018.

Joint main idea, derivations, numerical simulations, and write-up.

Introduction

This thesis consists of five papers concerned with the modeling and sampling of signals displaying structure in their spectral properties, i.e., in the distribution of power over frequency characterizing the signals. However, the specific meaning of "structure", as well as the motivation for studying it, differs among the papers. In the first two papers, it does not refer to the signal spectra themselves, but to the geometry of their domain of definition. Specifically, a notion of distance between signals that takes this geometry into account is proposed, and the implications of using this as a modeling tool is studied, yielding attractive robustness properties, as well as interpretability. In contrast, the structure considered in the last two papers refers to explicit parametric descriptions of the signals under study. Therein, it is studied how to optimally sample, i.e., to collect measurements from, a given signal as to allow for the best possible accuracy when estimating the signal parameters. Also, the implications of, deliberately, making simplified signal assumptions are studied. The third paper can be seen as a bridge between the first two and last two. Therein, it is studied how to view signals consisting of sinusoidal mixtures that are almost, but not quite, harmonically related. In this case, the structure arises both from the highly regular harmonic model, as well as in the geometry of the support. Using tools from the first two papers, a formal definition of the intuitive concept "almost" harmonic is proposed.

Below follows a brief introduction to the concepts and modeling tools used in the papers. Starting with outlining the connection between spectra and covariance sequences for stationary stochastic processes, the framework of optimal mass transport, which is the central modeling tool utilized in the first three papers, is introduced. After that, we revisit the Cramér-Rao lower bound as a tool for assessing achievable estimation performance, as well as the possibly less well-known misspecified version of the bound, setting the scene for the last three papers. A brief outline of the harmonic sinusoidal model is also provided.

1 Covariances and spectral representations

In applications of signal processing, the signal under study often displays some type of randomness or unpredictability, be it either due to actual stochasticity such as, e.g., thermal noise, or as a consequence of a simplified model of the physical process. A common strategy for handling this is to view the non-deterministic part of the signal as a stochastic process[1]. In this thesis, most signals under study fall in the category of wide-sense stationary (WSS) stochastic processes, and in particular, discrete-time, zero-mean processes in $\mathbb{C}$. That is, a signal $x(t) \in \mathbb{C}$, for $t \in \mathbb{Z}$, is said to be zero-mean WSS if $\mathbb{E}(x(t)) = 0$ for all t, and the covariance function

$$r_\tau \triangleq \mathbb{E}\left(x(t)\overline{x(t-\tau)}\right)$$

depends only on $\tau \in \mathbb{Z}$, for all t, where $\mathbb{E}(\cdot)$ and $\bar{\cdot}$ denote the expectation operation and complex conjugation, respectively. Such processes serve as models for many commonly encountered signals in, e.g., speech processing [1]. Furthermore, for any covariance sequence $\{r_k\}_{k\in\mathbb{Z}}$ of a WSS process, there exists a non-negative function Φ, referred to as the spectrum of the signal, defined on $\mathbb{T} \triangleq [-\pi, \pi)$, such that [2]

$$r_k = \frac{1}{2\pi}\int_{\mathbb{T}} \Phi(\omega)e^{ik\omega}d\omega, \tag{1}$$

for $k \in \mathbb{Z}$, where i is the imaginary unit. To identify Φ, i.e., the distribution of power over frequency, is often of central importance in applications, such as, e.g., noise suppression in speech enhancement tasks [3, 4], or for monitoring human diseases affecting the voice [5]. However, the spectral estimation problem, i.e., the task of finding a spectrum Φ from finite data lengths, is by no means a solved problem. Herein, finding Φ may be formulated as either

- estimate Φ from $x(t)$, for $t = 0, 1, \ldots, N-1$, for $N \in \mathbb{N}$, or
- estimate Φ from r_k, or estimates thereof, for $|k| \leq (n-1)$, for $n \in \mathbb{N}$.

It may be noted that for the latter formulation, spectral estimation may be viewed as an inverse problem, an interpretation that is not offered by the former. To elaborate on this, let $\mathcal{M}(\mathbb{T})$ denote the set of generalized integrable functions

[1] Herein, we will use the terms "signal" and "process" interchangeably.

on $\mathbb{T}$, and let $\mathcal{M}_+(\mathbb{T})$ denote the subset of $\mathcal{M}(\mathbb{T})$ consisting of non-negative functions. Thus, letting $\Phi \in \mathcal{M}_+(\mathbb{T})$ allows for spectra containing singular parts, i.e., Dirac delta functions, although integrals such as (1) then constitute some abuse of notation. Then, spectral estimation may be formulated as

- find $\Phi \in \mathcal{M}_+(\mathbb{T})$ such that (1) holds for $|k| \leq (n-1)$, for $n \in \mathbb{N}$.

However, the set of spectra consistent with the specified covariances, i.e.,

$$\Omega = \left\{ \Phi \in \mathcal{M}_+(\mathbb{T}) \mid r_k = \frac{1}{2\pi} \int_{\mathbb{T}} \Phi(\omega) e^{ik\omega} d\omega, \ |k| \leq (n-1) \right\},$$

often contains more than one element if n is finite, i.e., there is no unique spectrum Φ. In fact, the spectrum satisfying the covariance constraints is unique if and only if the Toeplitz covariance matrix defined by r_k, for, $k = 0, \pm 1, \ldots, \pm(n-1)$, is singular [2], i.e., if the rank of R is less than n, where

$$R = \begin{pmatrix} r_0 & r_{-1} & r_{-2} & \cdots & r_{-n+1} \\ r_1 & r_0 & r_{-1} & \cdots & r_{-n+2} \\ r_2 & r_1 & r_0 & \cdots & r_{-n+3} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ r_{n-1} & r_{n-2} & r_{n-3} & \cdots & r_0 \end{pmatrix}.$$

For the case when $\operatorname{rank}(R) = \tilde{n} < n$, the corresponding unique spectrum is of the form

$$\Phi(\omega) = 2\pi \sum_{k=1}^{\tilde{n}} \alpha_k^2 \delta(\omega - \omega_k),$$

where $\alpha_k \in \mathbb{R}_+$, $\omega_k \in \mathbb{T}$, and $\delta(\cdot)$ denotes the Dirac delta function, with the corresponding signal $x(t)$ being

$$x(t) = \sum_{k=1}^{\tilde{n}} \alpha_k e^{i\omega_k t + i\varphi_k},$$

where φ_k are uniform random variables on $\mathbb{T}$, independent of each other. Singular covariance matrices and their unique spectral representations are exploited in several classical line spectral estimation methods, such as Pisarenko [6], MUSIC [7],

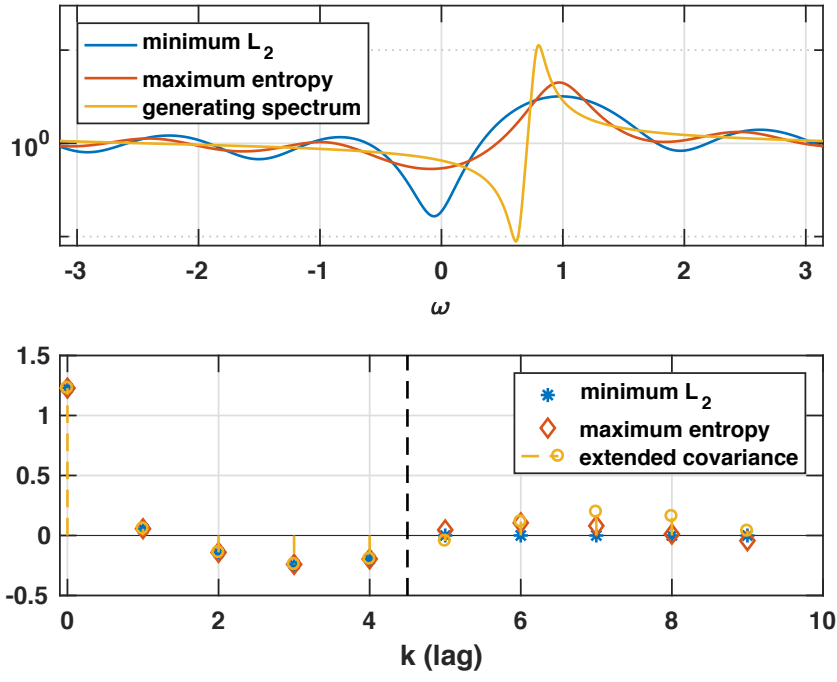

Figure 1: Top panel: three elements of Ω when covariances are specified up to lag $k = 4$ by a generating spectrum. Bottom panel: the real part of the covariance extensions, as implied by the three spectra.

ESPRIT [8], and, more recently, sparsity-aware techniques such as atomic norm minimization and denoising [9, 10]. However, if the covariance matrix is of rank n, which is the case for signals consisting of n or more sinusoidal components, or any signal with an absolutely continuous component[2], the spectrum is not uniquely defined, and additional assumptions are required for arriving at an estimate of Φ. It may be noted that the problem of uniqueness is not circumvented by considering estimating Φ directly from observations of the signal waveform. The non-uniqueness is illustrated in Figure 1, displaying a spectrum Φ used to generate r_k, for $k = 0, \pm 1, \ldots, \pm(n-1)$, with $n = 5$, in the top panel. Also shown are two other elements of Ω, i.e., spectra consistent with the observed covariances. Specifically, these two spectra are the elements of Ω with minimum L_2-norm and

[2]In this case, the covariance matrix is always full rank, irrespective of the dimension.

maximum entropy rate, respectively. Thus, depending on our assumptions, or preferences, one arrives at different, equally valid spectral estimates. However, different spectra will result in different evolutions of the covariance function, as shown in the bottom panel of Figure 1, displaying the real part of the covariance sequence extended up to lag $k = 9$ as implied by the three different spectra. As can be seen, the covariances differ after lag 4.

It may be noted that also the classical periodogram estimate, i.e.,

$$\Phi_{per}(\omega) = \frac{1}{N}\left|\sum_{t=0}^{N-1} x(t)e^{-i\omega t}\right|^2,$$

may be interpreted as a particular solution to the inverse covariance matching problem. Specifically, using the standard biased covariance estimates

$$\hat{r}_k = \frac{1}{N}\sum_{t=k}^{N-1} x(t)\overline{x(t-k)}, \tag{2}$$

for $k = 0, \pm 1, \ldots, \pm(N-1)$, the periodogram satisfies, through its connection to the correlogram [11, 12],

$$\frac{1}{2\pi}\int_{\mathbb{T}} \Phi_{per}(\omega)e^{ik\omega}d\omega = \begin{cases} \hat{r}_k, & \text{for } |k| \leq N-1 \\ 0, & \text{for } |k| > N-1. \end{cases}$$

That is, the periodogram specifies all covariances after a certain lag to be exactly zero, which is a quite strong assumption regarding the signal.

1.1 Spatial spectral estimation

Spectral estimation problems also appear in array processing applications such as localization and direction of arrival (DoA) estimation, in which one aims to determine the location of signal-emitting targets, or sources, based on measurements collected from a sensor array. The basis for doing this is to relate array measurements to a distribution of power over the space in which the sources are located. That is, letting Θ denote this space, one seeks $\Phi \in \mathcal{M}_+(\Theta)$ describing the distribution of signal power on Θ, and thereby identifying the locations of the sources, as $\Phi(\vartheta) > 0$ implies that signal is emitted from $\vartheta \in \Theta$. Here, one may have $\Theta \subset \mathbb{R}^d$ for the case of localization in d dimensional space, or $\Theta \subset \mathbb{T}$ for DoA

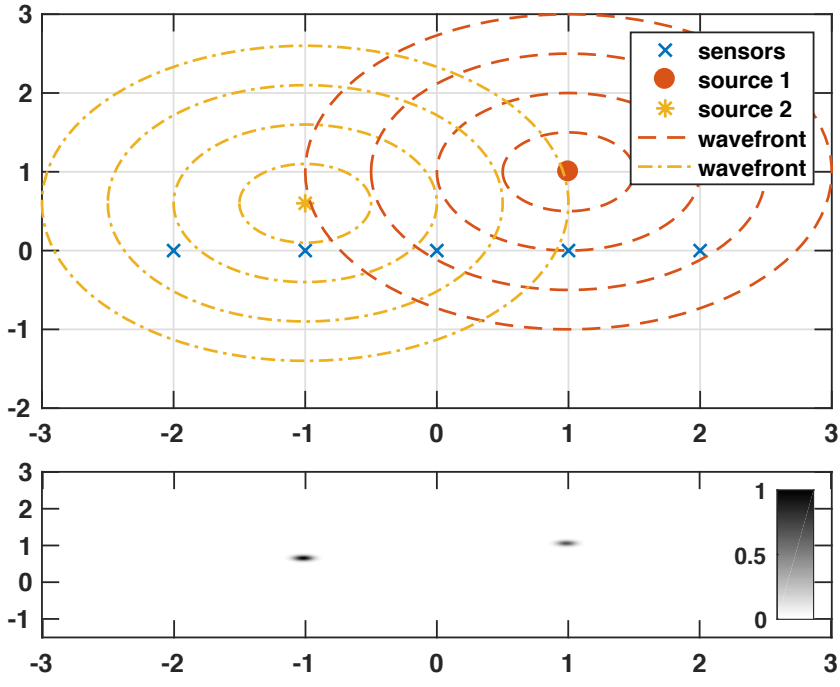

Figure 2: Top panel: Two signal sources impinging on an array consisting of five sensors. Bottom panel: corresponding spatial spectrum. The source power is indicated by the color intensity.

estimation. Spatial estimation problems are not limited to localizing objects; they also appear in, e.g., noise reduction applications [13, 14].

Let $y(t) \in \mathbb{C}^m$, for $t \in \mathbb{Z}$, denote a sequence of measurements corresponding to an array of $m \in \mathbb{N}$ sensors. Then, if the signals impinging on the array are WSS processes, so will $y(t)$ be. Furthermore, the covariance structure of $y(t)$, i.e., the sequence of so-called spatial covariance matrices

$$R(\tau) = \mathbb{E}\left(y(t)y(t-\tau)^H\right),$$

for $\tau \in \mathbb{Z}$, depends on the time-delays between targets and the sensors, and thus contains information of the target locations in Θ. If the impinging signals are sufficiently narrow-band, as well as independent of each other, this dependency allows for a simple expression relating $R = R(0)$ to Φ. Specifically, for narrow-band signals, the time-delay of a signal may be well-modeled as a phase-shift of

the corresponding carrier wave, i.e., there is a function $a : \Theta \to \mathbb{C}^m$, referred to as the array steering vector or manifold vector [15], such that $a(\vartheta)$ is the vector of phase-shifts, and potentially signal attenuations, corresponding to each of the array sensors for a signal impinging from $\vartheta \in \Theta$. It may here be noted that a is determined by the array geometry, the wave propagation properties of the space Θ, and the sensor response. Also, if the array geometry is unambiguous, i.e., a is injective on Θ, localization is possible in the sense that each $\vartheta \in \Theta$ is identifiable from $a(\vartheta)$. With this, the array covariance matrix R may be expressed as

$$R = \int_\Theta a(\vartheta)\Phi(\vartheta)a(\vartheta)^H d\vartheta. \tag{3}$$

In the case of finitely-many, say $K \in \mathbb{N}$, point-sources with signal powers α_k^2, for $k = 1, \ldots, K$, the array covariance takes on the commonly utilized form (see, e.g., [11])

$$R = \sum_{k=1}^{K} \alpha_k^2 a(\vartheta_k)a(\vartheta_k)^H = APA^H,$$

where

$$A = \begin{bmatrix} a(\vartheta_1) & \ldots & a(\vartheta_K) \end{bmatrix}, \; P = \operatorname{diag}\left(\begin{bmatrix} \alpha_1^2 & \ldots & \alpha_K^2 \end{bmatrix} \right),$$

with $\operatorname{diag}(x)$ denoting the diagonal matrix with x as its diagonal. As may be noted, the covariance matrix of a temporal process may be written on the form (3) by letting $\Theta = \mathbb{T}$, a be the Fourier vector of length m, and the measure of integration $d\vartheta = d\omega/2\pi$. The model in (3) will be used extensively in Paper B. An illustration of a spatial spectral estimation setup in two dimensions is displayed in Figure 2. Here, the top panel shows two signal sources emitting circular waves that impinge on an array consisting of five sensors. The bottom panel displays the corresponding spatial spectrum. As can be seen, the spectrum has two distinct peaks of power at the source locations. Note also that one of the modes is fainter than the other, corresponding to a lower power of the emitted signal.

The spatial spectral estimation problem, i.e., determining Φ from R, though being mathematically equivalent to the corresponding temporal problem, raises an interesting question that is not as apparent in the temporal case; how should distances between spectra, as well as their corresponding covariance matrices, be measured? In the spatial case, it is intuitively desirable that measures of spectral

similarity should reflect proximity on Θ, i.e., if two sources are closely located on Θ, in some native norm, then the corresponding spatial spectra ought also be "close". In particular, if a source at $\vartheta \in \Theta$ is perturbed to $\vartheta + \varepsilon$, where $\|\varepsilon\| \ll 1$, it seems reasonable that this should correspond to a small spectral perturbation. However, as we shall see next, such a property does not hold for some standard measures of similarity.

1.2 Spectral distances

As can be seen from Figure 1, spectra consistent with a certain number of covariances may differ significantly, which is also revealed by considering their respective full covariance sequences. To quantify such differences is of interest in estimation and classification applications. For example, the spectral representation of a sound recording may serve as the basis for identifying musical instruments or human speakers [16]. Also, in applications of spatial spectral estimation, the spectrum allows for locating and identifying objects [15]. To this end, one requires a means of comparing two spectra, i.e., determining a measure of distance between them. For example, for two sufficiently integrable spectra Φ_0, Φ_1, one could consider the L_2-norm

$$\|\Phi_0 - \Phi_1\|_2^2 = \int_{\mathbb{T}} (\Phi_0(\omega) - \Phi_1(\omega))^2 \, d\omega,$$

i.e., the integral of squared point-wise differences, or the Kullback-Leibler (KL) divergence [17],

$$D_{KL}(\Phi_0 \,||\, \Phi_1) = \int_{\mathbb{T}} \Phi_0(\omega) \log\left(\frac{\Phi_0(\omega)}{\Phi_1(\omega)}\right) d\omega,$$

i.e., the relative entropy between Φ_0 and Φ_1, commonly employed for measuring dissimilarity between probability distributions, and a popular choice as a fitness function in machine learning [18]. However, when used for quantifying the distance[3] between spectra Φ_0 and Φ_1, the comparison is mainly determined by the similarity of the spectra on the intersection of their supports. This is illustrated in Figure 3. As may be seen, the top panel displays three narrow-band spectra, all with the same shape, but with different center frequencies, each with a common flat noise-floor. For these spectra, the L_2-distances between the three resulting pairs are all the same as their supports are completely disjoint, save for the

[3] It may be noted that the KL divergence is not symmetric.

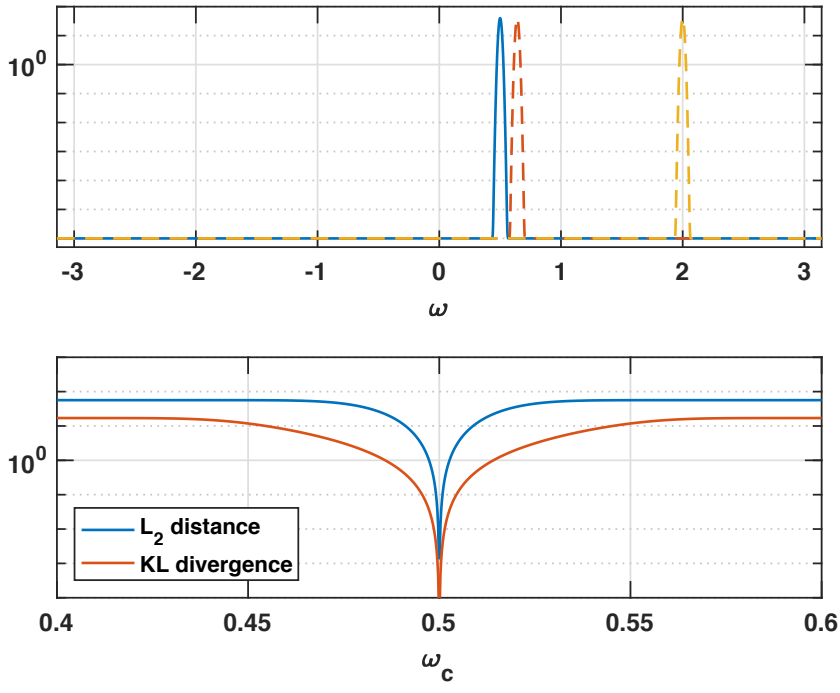

Figure 3: Top panel: three narrow-band spectra. Bottom panel: L_2-distance, as well as KL divergence, between spectrum with center frequency ω_c and spectrum with center frequency 0.5.

noise-floor. This is further illustrated in the bottom panel of Figure 3, displaying the L_2-distance between the spectrum with center frequency 0.5 and frequency-shifted copies of itself. As may be seen, the distance rapidly saturates as the center frequency is shifted. The corresponding effect is observed for the KL divergence. Thus, neither the L_2-distance[4] nor the KL-divergence are suitable measures of dissimilarity in applications in which the support of the spectra carry information. For example, in audio applications, the support of narrow-band spectra are related to the pitch perceived by human listeners, and thus, the proximity in frequency of the supports of two spectra affect the perceived similarity of the corresponding signals. In fact, the related effect of spectral masking is extensively used in audio coding [19]. As may be noted, the consequence of using distances not taking

[4]Indeed, any L_p-distance.

the structure of the support into account becomes more pronounced the more narrow-band the signals under consideration are. Furthermore, in spatial spectral estimation applications, such as DoA estimation and localization, the support of the spectrum corresponds directly to the location of the signal sources, i.e., the spectral support conveys information of physical proximity. Again, this is not reflected in neither the L_2-distance, nor the KL-divergence. Thus, in order to be meaningful for applications in which distances in frequency matter, and not only point-wise dissimilarity, the spectral distance measure needs to take the structure of the support into account explicitly. Such a distance may be obtained through the use of optimal mass transport (OMT), described next.

2 Optimal mass transport

As noted, a spectrum $\Phi_0 \in \mathcal{M}_+(\mathbb{T})$ describes the distribution of power over $\mathbb{T}$, and may thus be interpreted as a mass distribution. Then, in order to compare Φ_0 to some other $\Phi_1 \in \mathcal{M}_+(\mathbb{T})$, one may ask what the most efficient way of perturbing Φ_0 for it to become identical to Φ_1 is. Assigning a cost to this optimal perturbation then defines a notion of dissimilarity or distance between Φ_0 and Φ_1. This is the problem of OMT, as first formulated by Gaspard Monge (see, e.g., [20]). Specifically, the Monge problem of OMT seeks a mapping[5] $g : \mathbb{T} \to \mathbb{T}$ solving

$$\begin{aligned} \inf_{g} \quad & \int_{\mathbb{T}} c(\omega, g(\omega))\, \Phi_0(\omega) d\omega \\ \text{subject to} \quad & \int_{\omega_1 \in \mathcal{U}} \Phi_1(\omega_1) d\omega_1 = \int_{g(\omega_0) \in \mathcal{U}} \Phi_0(\omega_0) d\omega_0 \ , \ \text{for all } \mathcal{U} \subset \mathbb{T}, \end{aligned} \tag{4}$$

where $c : \mathbb{T} \times \mathbb{T} \to \mathbb{R}_+$ is a cost function. It may here be noted that as $c(\omega_0, \omega_1)$ describes the cost of moving one unit mass from ω_0 to ω_1, the value of the objective functional in (4) is equal to the total cost of moving all the mass in Φ_0 according to g. The constraint in (4), i.e., that the mass transported to any subset $\mathcal{U}$ is equal to the mass on the pre-image of $\mathcal{U}$, ensures that g indeed rearranges Φ_0 to Φ_1. It may be noted that, as a direct consequence, the total mass of Φ_0 and Φ_1 have to be equal,

$$\int_{\mathbb{T}} \Phi_0(\omega) d\omega = \int_{\mathbb{T}} \Phi_1(\omega) d\omega,$$

[5] We here adapt the formulation to the special case of mass distributions on $\mathbb{T}$.

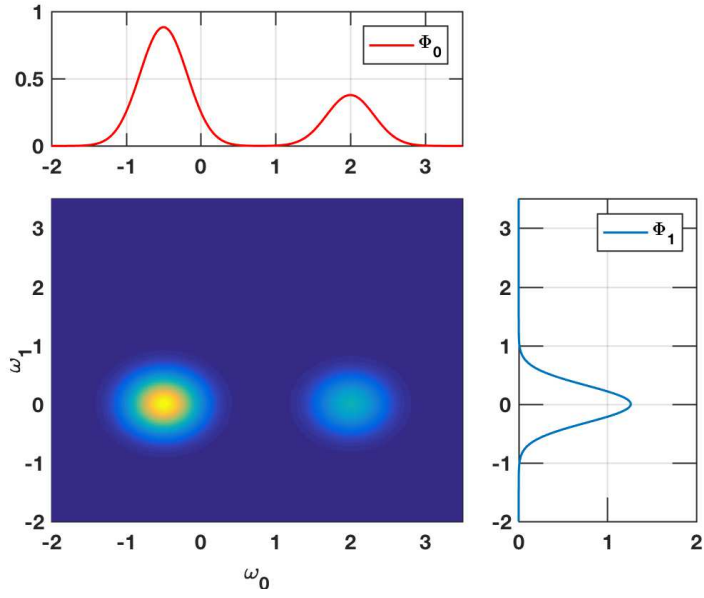

Figure 4: Two spectra, Φ_0, Φ_1, as well as a transport plan constructed as $M(\omega_0, \omega_1) = \Phi_0(\omega_0)\Phi(\omega_1)$.

i.e., they have to correspond to processes with the same variance. As for the cost function c, one may for example choose $c(\omega_0, \omega_1) = (\omega_0 - \omega_1)^2$, i.e., the cost of transporting a unit mass between two points is directly related to their distance on the frequency axis. In fact, for the case in which both Φ_0 and Φ_1 are absolutely continuous and when $c(\omega_0, \omega_1) = |\omega_0 - \omega_1|^p$ for some $p \geq 1$, the minimal value of (4) may be computed in closed form [21].

However, the minimum of (4) may not exist for general choices $\Phi_0, \Phi_1 \in \mathcal{M}_+(\mathbb{T})$. In fact, there might not even exist a g satisfying the constraint in (4). This is the case if Φ_0 and Φ_1 contain singular parts, as the formulation in (4), i.e., by modeling transport of mass using a mapping g, does not allow for splitting up point-masses. This renders the Monge formulation of OMT unsuitable for signal processing applications, as singular measures arising from sinusoidal components are common. Furthermore, even if the minimum of (4) exists, it is not easy to find the minimizing g, except for some special choices of the cost function c, due to the non-convexity of the problem. Both these problems may be remedied by,

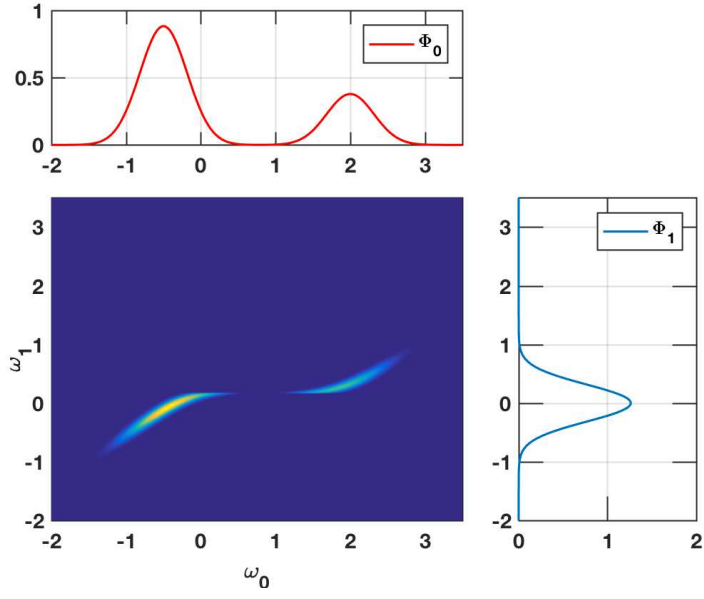

Figure 5: Two spectra, Φ_0, Φ_1, as well as a transport plan M resulting from solving (5) using $c(\omega_0, \omega_1) = (\omega_0 - \omega_1)^2$.

instead of using a mapping, modeling transport by means of a mass distribution on the product space $\mathbb{T} \times \mathbb{T}$. Specifically, one may let $M \in \mathcal{M}_+(\mathbb{T} \times \mathbb{T})$ describe transport between Φ_0 and Φ_1 by letting $M(\omega_0, \omega_1)$ be the mass transported from ω_0 to ω_1. This results in the Monge-Kantorovich problem of OMT, formulated as

$$\begin{aligned} \inf_{M \in \mathcal{M}_+(\mathbb{T}\times\mathbb{T})} \quad & \int_{\mathbb{T}} M(\omega_0, \omega_1) c(\omega_0, \omega_1) d\omega_0 d\omega_1 \\ \text{subject to} \quad & \int_{\mathbb{T}} M(\cdot, \omega) d\omega = \Phi_0 \ , \ \int_{\mathbb{T}} M(\omega, \cdot) d\omega = \Phi_1. \end{aligned} \tag{5}$$

Here, M is referred to a transport plan as it offers a complete description of the association of the masses of Φ_0 and Φ_1. Correspondingly, the so-called marginals of M must match Φ_0 and Φ_1, enforced by the constraints. It may here be noted that a feasible M always exists, as setting $M(\omega_0, \omega_1) = \Phi_0(\omega_0)\Phi(\omega_1)$, for all $\omega_0, \omega_1 \in \mathbb{T}$, results in a transport plan with the correct marginals. Furthermore,

in contrast to (4), the problem in (5) is a convex, even linear, program. Figures 4 and 5 provide an illustration of two feasible transport plans M for a pair of spectra Φ_0 and Φ_1, with Figure 4 displaying $M(\omega_0, \omega_1) = \Phi_0(\omega_0)\Phi(\omega_1)$ and Figure 5 showing the solution to (5), with $c(\omega_0, \omega_1) = (\omega_0 - \omega_1)^2$.

The use of OMT for modeling and estimation has recently attracted increasing attention in signal processing [21], machine learning [22], as well as in automatic control [23]. Furthermore, the framework has been used for endowing the space of spectra with a metric structure [24]. In Papers A, B, and C, we will see how tools based in OMT theory can be used for modeling distances, as well as transitions, between stochastic processes and signals.

3 Estimation of parametric models

Another type of structure considered in this thesis is that provided by parametric signal models, i.e., the signal is characterized by a parameter vector ϑ specifying its probability distribution. Herein, the considered ϑ are real and finite dimensional, i.e., elements of $\mathbb{R}^P$ for some $P \in \mathbb{N}$. As an example, consider the signal

$$x(t) = \alpha e^{-\beta t + i\omega t + i\varphi}, \tag{6}$$

where $\alpha > 0$, $\beta > 0$, $\omega \in [-\pi, \pi)$, and $\varphi \in [-\pi, \pi)$. This model is commonly employed in spectroscopy applications, wherein measurements are described as

$$y(t) = x(t) + e(t),$$

where $e(t)$ is a circularly symmetric white Gaussian noise with variance σ^2. For this signal, the parameter vector specifying the (Gaussian) probability density function (pdf) of a signal sample $\mathbf{y} \triangleq \begin{bmatrix} y(0) & y(1) & \dots & y(N-1) \end{bmatrix}^T$, for some $N \in \mathbb{N}$ is simply

$$\vartheta = \begin{bmatrix} \alpha & \beta & \omega & \varphi & \sigma^2 \end{bmatrix}^T.$$

It may be noted that the noise-free waveform $x(t)$ is completely deterministic. In Papers C, D, and, in particular, E, we are interested in assessing how well one may estimate the parameter ϑ from signal samples $\mathbf{y}$, i.e., what estimation variance may be achieved by an estimator $\hat{\vartheta}$. The main tool for doing this will be the Cramér-Rao lower bound (CRLB) (see, e.g., [25]), which we briefly outline here. Let

$p(\mathbf{y};\vartheta)$ denote the pdf of the signal sample[6] $\mathbf{y}$. If

$$\mathbb{E}\left(\nabla_\vartheta \log p(\mathbf{y};\vartheta)\right) = 0$$

for all ϑ, then,

$$\mathbb{E}\left((\hat{\vartheta}-\vartheta)(\hat{\vartheta}-\vartheta)^T\right) \succeq F(\vartheta)^{-1},$$

for any unbiased estimator $\hat{\vartheta}$, where $F(\vartheta)$ is the Fisher information matrix

$$F(\vartheta) = \mathbb{E}\left(\nabla_\vartheta \log p(\mathbf{y};\vartheta)\ \nabla_\vartheta \log p(\mathbf{y};\vartheta)^T\right), \tag{7}$$

and where $A \succeq B$ signifies that $A - B$ is positive semi-definite. Note here that all expectations are taken with respect to $p(\mathbf{y};\vartheta)$. Thus, $F(\vartheta)^{-1}$, referred to as the CRLB, is a lower bound, in the matrix sense, for the covariance matrix of estimators $\hat{\vartheta}$ and thus serves as benchmark for estimator performance. Furthermore, asymptotically, we expect the bound to be attained by the maximum likelihood estimator (MLE), i.e., the bound is informative. Asymptotically here means either as the sample length N or the signal to noise ratio (SNR), depending on application, tends to infinity.

3.1 Optimal sampling

The CRLB may not only be used as a benchmark for estimator variance; it can also serve as a tool for deciding how to optimally measure, or sample, a signal. In general, the FIM in (7) depends on where on its domain of definition the signal y has been sampled. For example, if $y(t)$ is a time series, the FIM depends on the sampling times. The interpretation of this is that different time instances, or combinations of sampling times, may carry different amount of information. The implication of this is that the CRLB will depend on the specific choice of samples collected, i.e., depending on which samples are selected, the difficulty of estimating the signal parameter ϑ will vary. In general, the more samples that are collected, the better. However, in some applications, there may be costs associated with collecting samples. For example, in Accordion nuclear magnetic resonance spectroscopy [26, 27], in which the signal consists of mixtures of decaying oscillating components such as (6), the sampling times t corresponds to a specific

[6] This is, for a general measurement of a signal, not necessarily corresponding to noisy measurements of (6).

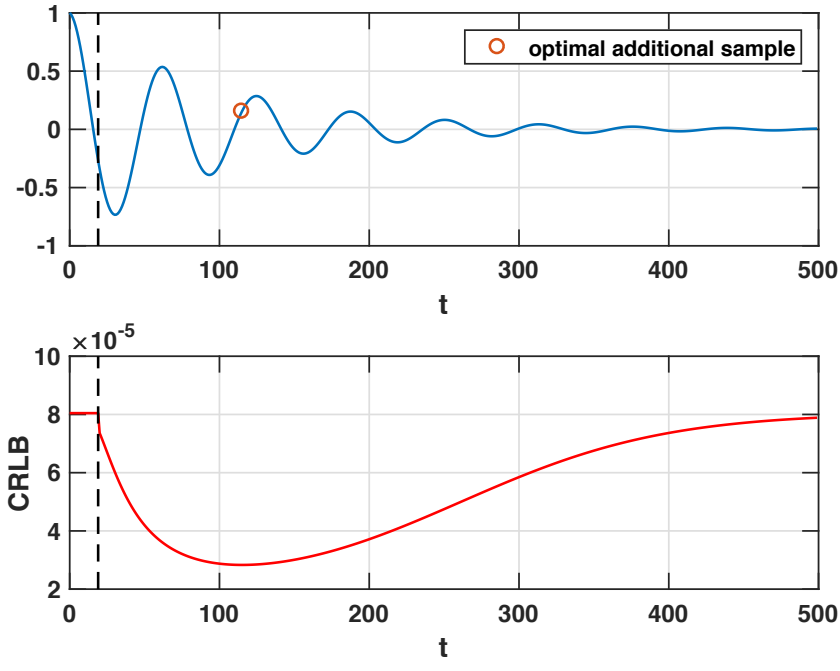

Figure 6: The problem of collecting one additional sample from (6) after having sampled at $t = 0, 1, \ldots, N_{initial} - 1$, with $N_{initial} = 20$. Top panel: real part of (noise-free) waveform. Bottom panel: CRLB for decay parameter β as a function of the additional sample collected.

experiment setting, with measurement $y(t) \in \mathbb{C}$ being the outcome of the experiment. Collecting signal samples is thus time consuming and one may pose the question: if only a certain number of signal measurements can be collected, how should these be selected as to allow for the best possible estimation performance?

As an illustration, consider having collected noisy samples from the model in (6) for $t = 0, 1, \ldots, N_{initial} - 1$, and then being offered to measure the signal at one additional time instance as to minimize the CRLB of β. Figure 6 displays such a situation, wherein $N_{initial} = 20$ samples have been collected. Here, the bottom panel displays the CRLB resulting from selecting one additional sample. As may be noted, the accuracy bound varies significantly depending on which sample is collected; the variance differs by almost a factor of 4 between the best and the worst sample. As may be noted, it is in this case most beneficial to sample

at $t = 115$, though it should be noted that this depends on the particular values of the signal parameters.

In Paper E, we address the general problem of sample selection for parametric models by formulating this, or rather a relaxed version of it, as a convex optimization problem. In particular, we utilize that, when samples collected at distinct points in the signal domain are statistically independent, the FIM decouples as a sum. For the temporal signal case, this means

$$F(\vartheta) = \sum_{t=0}^{N-1} F_t(\vartheta),$$

where

$$F_t(\vartheta) = \mathbb{E}\left(\nabla_\vartheta \log p(y(t);\vartheta)\ \nabla_\vartheta \log p(y(t);\vartheta)^T\right)$$

for $t = 0, 1, \ldots, N-1$, are information matrices corresponding to the individual samples. As may be noted, this property holds for deterministic signals observed in white noise.

3.2 Misspecified models

In some cases, the assumed signal model may not match that of the actual, measured signal. Specifically, letting $q(\cdot;\vartheta)$ denote the pdf of the assumed model, parametrized by the vector ϑ, and p denote the pdf of the measurements, the assumed model is said to be misspecified if there exists no ϑ such that $q(\cdot;\vartheta) \equiv p$. As an example, consider observing noisy measurements of (6) and, erroneously, assuming that the noise-free waveform is a pure sinusoid, i.e.,

$$z(t) = \alpha e^{i\omega t + i\varphi},$$

where, in this case, the parameter vector is $\vartheta = \begin{bmatrix} \alpha & \omega & \varphi & \sigma^2 \end{bmatrix}^T$. Clearly, for any decay parameter $\beta > 0$, this constitutes a misspecified model, as it is not possible to select the parameters α, ω, and φ as to match the waveform in (6), and thereby the pdfs cannot be made equal. However, one may still attempt to find a value of ϑ such that q is a good approximation of p. When the approximation is done in the Kullback-Leibler sense, this results in the concept of a pseudo-true parameter (see, e.g., [28]). Specifically, the pseudo-true parameter ϑ_0 is defined as

$$\vartheta_0 = \arg\min_\vartheta D_{KL}\left(p \,||\, q(\cdot;\vartheta)\right) = \arg\min_\vartheta \ -\mathbb{E}_p\left(\log q(\mathbf{y};\vartheta)\right),$$

where the expectation is taken with respect to p, i.e., the value ϑ_0 is such that the relative entropy between p and q is minimized. Furthermore, it can be shown that under some regularity conditions, the misspecified MLE, i.e., the MLE of ϑ derived under q, converges to ϑ_0 as the number of samples, or SNR, tends to infinity (see, [28, 29]). In addition, for estimators of ϑ_0 that are unbiased under p, i.e., satisfying

$$\mathbb{E}_p\left(\hat{\vartheta}_0\right) = \vartheta_0,$$

a lower bound on estimator variance may be found as the misspecified CRLB (MCRLB) [29], extending the standard CRLB. Thus, the concept of a pseudo-true parameter, together with the MCRLB, allows for analyzing the loss of performance, in terms of bias and estimator variance, when using misspecified models, for example due to deliberate model simplification, which is studied in Paper D. Furthermore, the misspecified modeling framework may serve as a basis for defining highly structured approximations of some signal classes. This is explored in Paper C, in which harmonic and inharmonic signals, briefly described next, are studied.

3.3 The harmonic model and inharmonicity

Consider the signal

$$x(t) = \sum_{k=1}^{K} \alpha_k e^{i\omega_k t + i\varphi_k}, \tag{8}$$

where $\alpha_k > 0$, $\varphi_k \in [-\pi, \pi)$, and $\omega_k = k\omega_0$, for $k = 1, 2, \ldots, K$, with $\omega_0 \in [-\pi, \pi)$. This model is referred to as harmonic, due to the integer relationship between the frequencies of the sinusoidal components and the fundamental frequency ω_0. The model in (8) is also referred to as a pitch and corresponds, in the context of audio signals, to the everyday concept of the same name. Harmonic signals appear in many fields of signal processing, ranging from the modeling of DNA sequences [30] to fault detection in industrial machinery [31]. However, it is most prominently featured in audio signal processing. For example, it has successfully been used for modeling the voiced part of human speech [32], as well as the sound produced by musical instruments [33]. Due to the ubiquity of the signal (8), a great number of methods for estimating its parameters from

noisy samples have been proposed, encompassing autocorrelation based methods exploiting the periodicity of $x(t)$ [34], as well as high-resolution sparse methods [35] (see also [36] for an overview). However, in some settings, the harmonic relationship $\omega_k = k\omega_0$ is only approximate. That is, the sinusoidal frequencies may be better described as

$$\omega_k = k\omega_0 + \Delta_k, \tag{9}$$

for $k = 1, \ldots, K$, where Δ_k are small perturbations. Such signals are often referred to as inharmonic [36]. In fact, there are complete parametric descriptions for some inharmonic signals, based on the physical properties of the generating mechanism. For example, the frequencies constituting the sound produced by vibrating strings are commonly modeled as [33]

$$\omega_k = \omega_0 k\sqrt{1 + \beta k^2},$$

where $\beta > 0$ is a parameter related to the string stiffness. There is also evidence for some degree of inharmonicity in the human voice, although no particular structure is known as of yet [37]. In the context of this thesis, the inharmonic model, i.e., the signal in (8) with frequencies (9) constitutes an example of a signal with both parametric and spectral structure. Specifically, letting φ_k be uniform random variables on $[-\pi, \pi)$, for $k = 1, 2, \ldots, K$, independent of each other, the signal in (9) is a WSS process with covariance function

$$r_\tau = \sum_{k=1}^{K} \alpha_k^2 e^{i\omega_k \tau}, \text{ for } \tau \in \mathbb{Z}.$$

Furthermore, the (unique) spectrum consistent with this covariance is

$$\Phi(\omega) = 2\pi \sum_{k=1}^{K} \alpha_k^2 \delta(\omega - \omega_k).$$

As may be noted, for the case of a perfectly harmonic signal, the spectrum consists of point-masses spaced evenly along the frequency axis by the fundamental frequency ω_0. In contrast, the spacing between the point-masses in the spectrum of an inharmonic signal deviates slightly from ω_0. An example is shown in Figure 7, displaying the periodogram spectral estimates of the musical note G_4 played by a

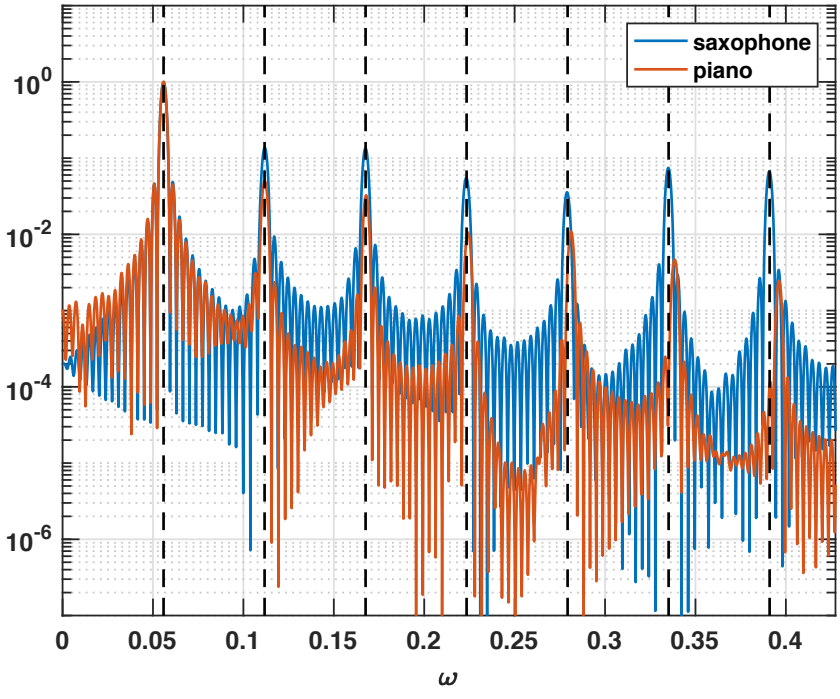

Figure 7: Periodogram spectral estimate of the musical note G_4 when played by a saxophone, as well as by a piano. The dashed lines indicate integer multiples of the nominal fundamental frequency.

saxophone, as well as a piano. The fundamental frequency of G_4 is 392 Hz, corresponding to an angular frequency of 0.0559 radians per sample when sampled at 44.1 kHz. As can be seen from the figure, the frequencies of the sinusoidal components of the saxophone signal correspond well to integer multiples of the fundamental frequency, whereas they deviate for the piano signal, i.e., the piano is inharmonic.

As may be noted, an inharmonic signal is not periodic, and thus lacks the concept of a fundamental frequency. Despite this, one may have the intuitive idea that an inharmonic signal is almost harmonic, although it is not clear in what sense. Paper C proposes a definition of fundamental frequency for non-periodic signals using tools from OMT as to formalize this intuitive idea.

4 Outline of the papers

Paper A: Interpolation and Extrapolation of Toeplitz Matrices via Optimal Mass Transport

In the first paper, the problem of quantifying dissimilarity between WSS stochastic processes based on their second-order statistics, i.e., covariance sequences or covariance matrices, is considered. Specifically, a notion of distance between Toeplitz covariance matrices is proposed, based on an optimal mass transport problem in the spectral domain. It is shown that the proposed distance has some attractive geometric properties, making it suitable for modeling perturbations of the frequency content of WSS processes. Furthermore, based on the proposed distance, a matrix interpolant is proposed, defining a path of valid covariance matrices connecting any two Toeplitz covariance matrices. Also, it is shown that, for certain choices of transport cost functions, the dual formulation of the OMT problem defining the distance allows for a finite-dimensional representation. The potential uses of the proposed distance and interpolant is illustrated in tracking and clustering examples. The work in Paper A has been published as

> Filip Elvander, Andreas Jakobsson, and Johan Karlsson, "Interpolation and Extrapolation of Toeplitz Matrices via Optimal Mass Transport". *IEEE Transactions on Signal Processing*, vol. 66, no. 20, pp. 5285-5298, October 2018.

Paper B: Multi-Marginal Optimal Transport using Partial Information

In the second paper, the idea of using OMT as a means of measuring distances between objects parametrized by mass distributions, such as covariance matrices parametrized by spectra, is generalized to modeling connections and proximity between collections of multiple objects. This is referred to as multi-marginal problems with partial information. It is shown that problems in spatial spectral estimation and array processing, such as sensor fusion and tracking, may be cast in this form. Furthermore, a main contribution of this work is proposing a computationally efficient algorithm capable of addressing high-dimensional transport problems. The properties of the proposed modeling framework, as well as the computational tools, are demonstrated in numerical examples from array processing, such as tracking and sensor fusion, wherein the sensor arrays are not jointly calibrated. It is shown that the OMT formulation induces inherent ro-

bustness to modeling errors, as well as allows for incorporating knowledge of underlying dynamics governing changes in the spectral properties of the observed signals. The work in Paper B has been published as

> Filip Elvander, Isabel Haasler, Andreas Jakobsson, and Johan Karlsson, "Multi-marginal optimal transport using partial information with applications in robust localization and sensor fusion". *Signal Processing*, vol. 171, pp. 1-19, 2020.

Paper C: Defining Fundamental Frequency for Almost Harmonic Signals

The third paper studies close-to-harmonic signals, and proposes three definitions of a concept of fundamental frequency, or pitch, for inharmonic, i.e., non-periodic, signals, emanating from three different approximation frameworks and interpretations of the signal. In all three definitions, the inharmonic signal is associated with a harmonic counterpart whose fundamental frequency then serves as the definition of pitch for the inharmonic signal. The first definition views the harmonic companion signal as a misparametrization of the inharmonic counterpart and uses tools from the framework of misspecified estimation in defining the fundamental frequency. The second definition instead considers the spectral properties of the inharmonic signal, viewing it as a spectral perturbation of a harmonic counterpart, and uses OMT theory for formally defining the fundamental frequency. In contrast to the first and second definitions, the third interprets inharmonicity as a stochastic frequency perturbation, and defines the fundamental frequency via an expectation. The properties of the three definitions, as well as the consequences of using either of them in estimation, is theoretically explored, as well as illustrated in simulation studies. The work in Paper C has been submitted for possible publication as

> Filip Elvander and Andreas Jakobsson, "Defining Fundamental Frequency for Almost Harmonic Signals".

Paper D: Mismatched Estimation of Polynomially Damped Signals

In the fourth paper, the implications of using simplified models in estimation of the parameters of spectroscopic signals is studied. Specifically, posing the problem within the framework of misspecified estimation, the so-called pseudo-true parameters defining the reduced signal models are used for predicting the behavior of

estimation methods operating under simplified assumptions. It is shown that the higher-order terms in spectroscopic signals may estimated sequentially, allowing for reduced computational load. The work in Paper D has been published in part as

> Filip Elvander, Johan Swärd, and Andreas Jakobsson, "Mismatched Estimation of Polynomially Damped Signals". *Proceedings of the 8th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, pp. 246-250, 2019.

Paper E: Designing Sampling Schemes for Multi-Dimensional Data

The fifth and final paper considers the problem of optimally sampling, or collecting measurements from, structured signals as to allow for the best possible estimation accuracy, as measured by the Cramér-Rao lower bound. Posing this as a selection problem, a convex relaxation in the form of a semi-definite program is proposed for approximating the original combinatorial problem. The resulting formulation allows for selecting arbitrary subsets, or positive linear combinations, of the signal parameters, thus enabling concentrating on the parameters of interest and ignoring those considered as nuisance. The proposed method also allows for modeling imprecise knowledge of the signal parameters. As shown in numerical examples, the method allows for a dramatic reduction in the number of signal samples required for achieving a given bound on estimation accuracy. The tools of Paper E have recently been applied to nuclear magnetic resonance spectroscopy experiments, allowing for more than a ten-fold reduction in the required experiment time [27]. The work in Paper E has been published as

> Johan Swärd, Filip Elvander, and Andreas Jakobsson, "Designing sampling schemes for multi-dimensional data". *Signal Processing*, vol. 150, pp. 1-10, 2018.

References

[1] P. Kroon and W. B. Kleijn, "Linear-prediction based analysis-by-synthesis coding," in *Speech Coding and Synthesis*, W. B. Kleijn and K. K. Paliwal, Eds., chapter 3, pp. 79–119. Elsevier, Berlin, Germany, 1995.

[2] U. Grenander and G. Szegö, *Toeplitz Forms and Their Applications*, University of California Press, Los Angeles, 1958.

[3] J. Benesty, J. R. Jensen, M. G. Christensen, and J. Chen, *Speech Enhancement: A Signal Subspace Perspective*, Academic Press, 2014.

[4] J. R. Jensen, J. Benesty, M. G. Christensen, and S. H. Jensen, "Enhancement of single-channel periodic signals in the time-domain," *IEEE Trans. Audio, Speech, Language Process*, vol. 20, no. 7, pp. 1948–1963, 2012.

[5] M. A. Little, P. E. McSharry, E. J. Hunter, J. Spielman, and L. O. Ramig, "Suitability of Dysphonia Measurements for Telemonitoring of Parkinson's disease," *IEEE Trans. Biomed. Eng.*, vol. 56, no. 4, pp. 1015–102, April 2009.

[6] V. F. Pisarenko, "The retrieval of harmonics from a covariance function," *Geophys. J. Royal Astronomical Soc.*, vol. 33, pp. 347–366, 1973.

[7] R. Schmidt, "Multiple emitter location and signal parameter estimation," in *Proceedings of RADC Spectrum Estimation Workshop*, 1979, pp. 243–258.

[8] A. Paulraj, R. Roy, and T. Kailath, "A Subspace Rotation Approach to Signal Parameter Estimation," *Proceedings of the IEEE*, vol. 74, no. 7, pp. 1044–1045, July 1986.

[9] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky, "The Convex Geometry of Linear Inverse Problems," *Foundations of Computational Mathematics*, vol. 12, no. 6, pp. 805–849, Dec 2012.

[10] B. N. Bhaskar, G. Tang, and B. Recht, "Atomic Norm Denoising with Applications to Line Spectral Estimation," *IEEE Trans. Signal Process.*, vol. 61, no. 23, pp. 5987 – 5999, July 2013.

[11] P. Stoica and R. Moses, *Spectral Analysis of Signals*, Prentice Hall, Upper Saddle River, N.J., 2005.

[12] R. B. Blackman and J. W. Tuckey, *The Measurement of Power Spectra from the Point of View of Communications Engineering*, Dover, New York, 1959.

[13] F. Elvander, R. Ali, A. Jakobsson, and T. van Waterschoot, "Offline Noise Reduction Using Optimal Mass Transport Induced Covariance Interpolation," in *27th European Signal Processing Conference*, A Coruna, Spain, Sept. 2019.

[14] J. Zhang, S. P. Chepuri, R. C. Hendriks, and R. Heusdens, "Microphone subset selection for MVDR beamforming based noise reduction," *IEEE/ACM Trans. on Audio, Speech, and Language Processing*, vol. 26, no. 3, pp. 550–563, 2018.

[15] H. L. Van Trees, *Detection, Estimation, and Modulation Theory: Radar-Sonar Signal Processing and Gaussian Signals in Noise*, Krieger Publishing Co., Inc., 1992.

[16] A. Klapuri and M. Davy, *Signal Processing Methods for Music Transcription*, Springer, 2006.

[17] S. Kullback and R. A. Leibler, "On Information and Sufficiency," *The Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79–86, March 1951.

[18] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," in *2nd International Conference on Learning Representations*, Banff, Canada, 2014.

[19] T. Painter and A. S. Spanias, "Perceptual coding of digital audio," *Proceedings of IEEE*, vol. 88, no. 4, pp. 451–515, April 2000.

[20] C. Villani, *Optimal transport: old and new*, Springer Science & Business Media, 2008.

[21] S. Kolouri, S. R. Park, M. Thorpe, D. Slepcev, and G. K. Rohde, "Optimal Mass Transport: Signal processing and machine-learning applications," *IEEE Signal Process. Mag.*, vol. 34, no. 4, pp. 43–59, July 2017.

[22] M. A. Schmitz, M. Heitz, N. Bonneel, F. Ngolè, D. Coeurjolly, M. Cuturi, G. Peyré, and J-L. Starck, "Wasserstein Dictionary Learning: Optimal Transport-Based Unsupervised Nonlinear Dictionary Learning," *SIAM J. Imaging Sci.*, vol. 11, no. 1, pp. 643–678, 2018.

[23] Y. Chen, T. T. Georgiou, and M. Pavon, "Optimal Transport Over a Linear Dynamical System," *IEEE Trans. Autom. Control*, vol. 62, no. 5, pp. 2137–2152, May 2017.

[24] T. T. Georgiou, J. Karlsson, and M. S. Takyar, "Metrics for power spectra: an axiomatic approach," *IEEE Trans. Signal Process.*, vol. 57, no. 3, pp. 859–867, Mar. 2009.

[25] S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory*, Prentice-Hall, Englewood Cliffs, N.J., 1993.

[26] A. M. Mandel and A. G. Palmer III, "Measurement of Relaxation-Rate Constants Using Constant-Time Accordion NMR Spectroscopy," *J. Magn. Reson., Series A*, vol. 110, pp. 62–72, 1994.

[27] G. Carlström, F. Elvander, J. Swärd, A. Jakobsson, and M. Akke, "Rapid NMR Relaxation Measurements Using Optimal Non-Uniform Sampling of Multi-Dimensional Accordion Data Analyzed by a Sparse Reconstruction Method," *J. Phys. Chem. A*, 2019.

[28] S. Fortunati, F. Gini, M. S. Greco, and C. D. Richmond, "Performance bounds for parameter estimation under misspecified models: Fundamental findings and applications," *IEEE Signal Process. Mag.*, vol. 34, no. 6, pp. 142–157, Nov 2017.

[29] C. D. Richmond and L. L. Horowitz, "Parameter Bounds on Estimation Accuracy Under Model Misspecification," *IEEE Trans. Signal. Process*, vol. 63, no. 9, pp. 2263–2278, 2015.

[30] L. Galleani and R. Garello, "The Minimum Entropy Mapping Spectrum of a DNA Sequence," *IEEE Trans. Inf. Theor.*, vol. 56, no. 2, pp. 771–783, Feb. 2010.

[31] R. B. Randall, *Vibration-Based Condition Monitoring: Industrial, Aerospace and Automotive Applications*, John Wiley & Sons, Chichester, UK, 2011.

[32] S. M. Nørholm, J. R. Jensen, and M. G. Christensen, "Instantaneous Fundamental Frequency Estimation With Optimal Segmentation for Non-stationary Voiced Speech," *IEEE/ACM Trans. Audio, Speech, Language Process.*, vol. 24, no. 12, pp. 2354–2367, Dec 2016.

[33] N. H. Fletcher and T. D. Rossing, *The Physics of Musical Instruments*, Springer-Verlag, New York, NY, 1988.

[34] A. de Cheveigné and H. Kawahara, "YIN, a fundamental frequency estimator for speech and music," *J. Acoust. Soc. Amer.*, vol. 111, no. 4, pp. 1917–1930, April 2002.

[35] T. L. Jensen and L. Vandenberghe, "Multi-pitch Estimation using Semidefinite Programming," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, New Orleans, LA, USA, March 2017.

[36] M. Christensen and A. Jakobsson, *Multi-Pitch Estimation*, Morgan & Claypool, San Rafael, Calif., 2009.

[37] E. B. George and M. J. T. Smith, "Speech analysis/synthesis and modification using an analysis-by-synthesis/overlap-add sinusoidal model," *IEEE Trans. Speech Audio Process.*, vol. 5, no. 5, pp. 389–406, Sep 1997.

A

Paper A

Interpolation and Extrapolation of Toeplitz Matrices via Optimal Mass Transport

Filip Elvander[1], Andreas Jakobsson[1], and Johan Karlsson[2]

[1] *Centre for Mathematical Sciences, Lund University, Lund, Sweden*
[2] *Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden*

Abstract

In this work, we propose a novel method for quantifying distances between Toeplitz structured covariance matrices. By exploiting the spectral representation of Toeplitz matrices, the proposed distance measure is defined based on an optimal mass transport problem in the spectral domain. This may then be interpreted in the covariance domain, suggesting a natural way of interpolating and extrapolating Toeplitz matrices, such that the positive semi-definiteness and the Toeplitz structure of these matrices are preserved. The proposed distance measure is also shown to be contractive with respect to both additive and multiplicative noise, and thereby allows for a quantification of the decreased distance between signals when these are corrupted by noise. Finally, we illustrate how this approach can be used for several applications in signal processing. In particular, we consider interpolation and extrapolation of Toeplitz matrices, as well as clustering problems and tracking of slowly varying stochastic processes.

Key words: Covariance interpolation, optimal mass transport, Toeplitz matrices, spectral estimation

1 Introduction

Statistical modeling is a key methodology for estimation and identification and is used throughout the signal processing field. An intrinsic component of such models is covariance estimates, which are extensively used in application areas such as spectral estimation, radar, and sonar [1, 2], wireless channel estimation, medical imaging, and identification of systems and network structures [3, 4]. Although being a classical subject (see, e.g., [5]), covariance estimation has recently received considerable attention. Such contributions include works on finding robust covariance estimates with respect to outliers, as well as methods suitable for handling different distribution assumptions, including families of non-Gaussian distributions [6–10]. Another important class of problems is covariance estimation with an inherent geometry that gives rise to a structured covariance matrix. Such structures could arise from stationarity assumptions of the underlying object [11–15] or be due to assumptions in, e.g., the underlying network structures in graphical models [16, 17]. In this work, we focus on Toeplitz structures which naturally arise when modeling stationary signals and processes.

Although many methods rely on stationarity for modeling signals, such assumptions are typically not valid over longer time horizons. Therefore, tools for interpolation and morphing of covariance matrices are important for modeling and fusing information. A straightforward and often used approach is the Euclidean metric; however, this metric does not take into account the underlying geometry and typically results in fade-in-fade-out effects (as is also illustrated herein). Several other such tools for interpolating covariances have recently been proposed in the literature, for example methods based on g-convexity [9], optimal mass transport [18], and information geometry [19]. An alternative approach for such interpolation is to relax, or "lift", the covariances and instead consider interpolation between the lifted objects. For example, in [20] (see also [21]), interpolation between covariance matrices is induced from the optimal mass transport geodesics between the Gaussian density functions with the corresponding covariances. However, neither of these interpolation approaches take into account that the covariance matrix represents an (almost) stationary times series and do not preserve the Toeplitz structure of the interpolating covariance sequence.

The topic of optimal mass transport (see, e.g., [22, 23]) was originally introduced in order to address the problem of, in a cost efficient way, supplying construction sites with building material and has been used in many contexts, such as, e.g., economics and resource allocation. Lately, it has also gained interest

in application fields such as image processing [24, 25] signal processing [26–28], computer vision and machine learning [29–33]. In this work, we will utilize optimal mass transport to model changes in the covariance structure of stochastic processes, or signals. To this end, we propose a new lifting approach, where the lifting is made from the covariance domain to the frequency domain, using the fact that any positive semi-definite Toeplitz matrix has a spectral representation. We combine this approach with the frequency domain metric based on optimal mass transport, proposed in [27], in order to define pairwise distances between Toeplitz matrices. This is done by considering the minimum distance, in the optimal mass transport sense, between the sets of power spectra consistent with each of the Toeplitz matrices. The proposed distance measure is shown to be contractive with respect to additive and multiplicative noise, i.e., it reflects the increased difficulty of discriminating between two stochastic processes if these are corrupted by two realizations of a noise process. Also, we show that the proposed distance measure gives rise to a natural way of interpolating and extrapolating Toeplitz matrices. The interpolation method preserves the Toeplitz structure, the positive semi-definiteness, as well as the diagonal of the interpolating/extrapolating matrices.

The proposed optimal mass transport problem is in its original form an infinite-dimensional problem. As an alternative to finding solutions using approximations based on discretizations of the underlying space, we show that certain formulations of the problem allows for approximations by a semi-definite program using a sum-of-squares representation. Also, we illustrate how the method can be used for interpolation, extrapolation, tracking, and clustering.

This paper is organized as follows. In Section 2, we provide a brief background on the moment problem, i.e., determining the power spectrum from a partial covariance sequence or finite covariance matrix, as well as introduce the problem of optimal mass transport. Section 3 introduces the proposed distance notion for positive semi-definite Toeplitz covariance matrices. Here, the dual problem is derived, and properties of the proposed distance notion are described. In Section 4, we describe applications of the proposed distance notion, such as induced interpolation, extrapolation, tracking, and clustering. Section 5 formulates a sum-of-squares relaxation of the dual problem. Section 6 provides numerical illustrations of the proposed distance notion, as well as the described applications. Finally, Section 7 concludes upon the work.

Notation

Let $\mathbb{M}^n$ denote the set of Hermitian $n \times n$ matrices and let $(\cdot)^T$ denote the transpose, $(\cdot)^H$ the Hermitian transpose, and $\overline{(\cdot)}$ the complex conjugate. Let $\mathbb{T} = [-\pi, \pi)$ and let $C_{\text{perio}}(\mathbb{T})$ denote the set of continuous and 2π-periodic functions on $\mathbb{T}$. The set of linear bounded functionals on $C_{\text{perio}}(\mathbb{T})$, or equivalently, the dual space of $C_{\text{perio}}(\mathbb{T})$, is the set of generalized integrable functions on the set $\mathbb{T}$, here denoted by $\mathcal{M}(\mathbb{T})$. Thus, $\mathcal{M}(\mathbb{T})$ includes, e.g., functions containing singular parts such as Dirac delta functions [34].[1] Further, we let $\mathcal{M}_+(\mathbb{T})$ denote the subset of such functions that are non-negative. We use $\langle \Phi, f \rangle$ to denote the application of the functional Φ on f, e.g.,

$$\langle \Phi, f \rangle = \int_{\mathbb{T}} f(\vartheta) \Phi(\vartheta) d\vartheta$$

if $f \in C_{\text{perio}}(\mathbb{T})$ and $\Phi \in \mathcal{M}(\mathbb{T})$. For Hilbert spaces, $\langle \cdot, \cdot \rangle$ is the standard inner product, e.g., when $\mathbf{X}$ and $\mathbf{Y}$ are vectors or matrices, then $\langle \mathbf{X}, \mathbf{Y} \rangle = \text{tr}(\mathbf{X}\mathbf{Y}^H)$, where $\text{tr}(\cdot)$ denotes the trace. We denote matrices by boldface upper-case letters, such as $\mathbf{X}$, whereas vectors are denoted by boldface lower-case letters, such as $\mathbf{x}$. Furthermore, $\|\mathbf{X}\|_F = \sqrt{\langle \mathbf{X}, \mathbf{X} \rangle}$ denotes the Frobenius norm induced by the matrix inner product. Lastly, for $f \in C_{\text{perio}}(\mathbb{T})$, we let

$$\|f\|_1 = \sup_{|\Phi(\vartheta)| \leq 1} \langle \Phi, f \rangle = \int_{\mathbb{T}} |f(\vartheta)| \, d\vartheta$$

denote the L_1-norm.

2 Background

2.1 Stochastic processes and spectral representations

We will in this work consider complex-valued discrete time stochastic processes, or signals, $y(t)$ for $t \in \mathbb{Z}$. These will be assumed to be zero mean and wide sense stationary (WSS), i.e., $\mathbb{E}(y(t)) = 0$ for all $t \in \mathbb{Z}$, and the covariance

$$r_k \triangleq \mathbb{E}(y(t)\overline{y(t-k)}) \tag{1}$$

[1] Strictly speaking $\mathcal{M}(\mathbb{T})$ is the set of signed bounded measures on $\mathbb{T}$ [35].

being independent of t. Here, $\mathbb{E}(\cdot)$ denotes the expectation operator. The frequency content of the process $y(t)$ may then be represented by the power spectrum, Φ, i.e., the non-negative function on $\mathbb{T}$ whose Fourier coefficients coincide with the covariances:

$$r_k = \frac{1}{2\pi}\int_{-\pi}^{\pi} \Phi(\vartheta) e^{-ik\vartheta} d\vartheta \tag{2}$$

for $k \in \mathbb{Z}$ (see, e.g., [36, Chapter 2]). Typically in spectral estimation, one considers the inverse problem of recovering the power spectrum Φ from a given set of covariances r_k, for $k \in \mathbb{Z}$, with $|k| \leq n-1$. The condition for any such reconstruction to be valid is that Φ should be consistent with the covariance sequence $\{r_k\}_{|k|\leq n-1}$, i.e., (2) should hold for $|k| \leq n-1$. The corresponding $n \times n$ covariance matrix is

$$\mathbf{R} = \begin{pmatrix} r_0 & r_{-1} & r_{-2} & \cdots & r_{-n+1} \\ r_1 & r_0 & r_{-1} & \cdots & r_{-n+2} \\ r_2 & r_1 & r_0 & \cdots & r_{-n+3} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ r_{n-1} & r_{n-2} & r_{n-3} & \cdots & r_0 \end{pmatrix} \tag{3}$$

which is a Hermitian Toeplitz matrix, since $y(t)$ is WSS. Thus, expressed in the form of matrices, a spectrum is consistent with an observed partial covariance sequence, or, equivalently, a finite covariance matrix, if $\Gamma(\Phi) = \mathbf{R}$, where $\Gamma : \mathcal{M}(\mathbb{T}) \to \mathbb{M}^n$ is the linear operator

$$\Gamma(\Phi) \triangleq \frac{1}{2\pi}\int_{\mathbb{T}} \mathbf{a}(\vartheta)\Phi(\vartheta)\mathbf{a}(\vartheta)^H d\vartheta \tag{4}$$

and

$$\mathbf{a}(\vartheta) \triangleq \begin{bmatrix} 1 & e^{-i\vartheta} & \cdots & e^{-i(n-1)\vartheta} \end{bmatrix}^T \tag{5}$$

is the Fourier vector. Note that $\Gamma(\Phi)$ is a Toeplitz matrix, since $\mathbf{a}(\vartheta)\mathbf{a}(\vartheta)^H$ is Toeplitz for any ϑ. It may be noted that for any positive semi-definite Toeplitz matrix, $\mathbf{R}$, there always exists at least one consistent power spectrum; in fact, if $\mathbf{R}$ is positive definite, there is an infinite family of consistent power spectra [37]. It may be noted that for singular Toeplitz covariance matrices, the spectral representation is unique. This fact has recently been successfully utilized in atomic

norm minimization problems for grid-less compressed sensing of sinusoidal signals (see, e.g., [38,39]). In this work, we are mainly interested in the non-singular case, where several power spectra are consistent with given covariance matrices. In Section 3, we will utilize such spectral representations in order to define a notion of distance between pairs of Toeplitz matrices. This distance will be defined in terms of the minimum optimal mass transport cost between the sets of power spectra consistent with the matrices.

2.2 Optimal mass transport

The Monge-Kantorovich problem of optimal mass transport is the problem of finding an optimal transport plan between two given mass distributions [22,23]. The cost of moving a unit mass is defined on the underlying space, and the optimal transport plan is defined as the plan minimizing the total cost. The resulting minimal cost, associated with the optimal transport plan, can then be used as a measure of similarity, or distance, between the two mass distributions. The idea of utilizing the optimal mass transport cost as a distance measure has been used, e.g., for defining metrics on the space of power spectra [27], whereas the optimal transport plan has been used for tracking stochastic processes with smoothly varying spectral content and for spectral morphing for speech signals [28]. Recently, the interpretation of the optimal transport plan as providing an optimal association between elements in two mass distributions has been used as a means of clustering in fundamental frequency estimation algorithms [40]. One of the advantages of using the optimal mass transport as a distance compared to traditional metrics is that it naturally incorporates the geometry of the underlying space. In particular, the optimal mass transport cost between two objects depends on the distance between the two objects in the underlying space, whereas standard metrics only depend on the overlapping regions. Furthermore, interpolation using optimal mass transport results in smooth transitions in the underlying space (see Sections 6.2 and 6.3). This makes optimal mass transport suitable for applications where there is a smooth transition in the underlying space, e.g., tracking problems in radar and sonar.

As in [27], we consider the following distance between two spectra Φ_0 and Φ_1:

$$S(\Phi_0, \Phi_1) \triangleq \min_{M \in \mathcal{M}_+(\mathbb{T}^2)} \int_{\mathbb{T}^2} c(\vartheta, \varphi) M(\vartheta, \varphi) d\vartheta d\varphi \tag{6a}$$

$$\text{subject to} \quad \Phi_0(\vartheta) = \int_{\mathbb{T}} M(\vartheta, \varphi) d\varphi \tag{6b}$$

$$\Phi_1(\varphi) = \int_{\mathbb{T}} M(\vartheta, \varphi) d\vartheta, \tag{6c}$$

where $\mathbb{T}^2 = \mathbb{T} \times \mathbb{T}$ denotes the 2-D frequency space. Here, the cost function, $c(\vartheta, \varphi)$, details the cost of moving one unit of mass between the frequencies ϑ and φ. The transport plan, $M(\vartheta, \varphi)$, specifies the amount of mass moved from frequency ϑ to frequency φ. The objective in (6a) is the total cost associated with the transport plan M and the constraints (6b) and (6c) ensure that M is a valid transport plan from Φ_0 to Φ_1, i.e., the integration marginals of M coincide with the spectra Φ_0 and Φ_1. It may be noted that, due to these marginal constraints, the distance measure S is only defined for spectra of the same mass, or total power. However, S may be generalized in order to allow for mass differences by including a cost for adding and subtracting mass by postulating that the spectra Φ_0 and Φ_1 are perturbations of functions Ψ_0 and Ψ_1 that have equal mass. As in [27], this may be formulated as

$$S_\varkappa(\Phi_0, \Phi_1) = \min_{\Psi_j \in \mathcal{M}_+(\mathbb{T})} S(\Psi_0, \Psi_1) + \varkappa \sum_{j=0}^{1} \left\| \Phi_j - \Psi_j \right\|_1 \tag{7}$$

where $\varkappa > 0$ is a used-defined parameter detailing the cost of adding or subtracting mass. One interpretation of this is that points that are close represent the same object and can thus be transported via the first term in (7), whereas points that are far apart represent different objects and must be phased in/out using the second term in (7). Then, $\varkappa$ may be interpreted as a parameter determining when two points are close. If the cost function in the optimal mass transport problem in (6) is chosen as $c(\vartheta, \varphi) = d(\vartheta, \varphi)^p$, $p \geq 1$ for any metric $d(\vartheta, \varphi)$ on $\mathbb{T}$, then $W(\Phi_0, \Phi_1) = S(\Phi_0, \Phi_1)^{1/p}$ is the so-called Wasserstein metric on $\mathcal{M}_+(\mathbb{T})$. Similarly, for $S_\varkappa$, the following theorem holds.

Theorem 1 ([27]). *Let $p \geq 1$, $\varkappa > 0$, and let the cost function be $c(\vartheta, \varphi) = |\vartheta - \varphi|^p$. Then,*

$$W_\varkappa(\Phi_0, \Phi_1) = S_\varkappa(\Phi_0, \Phi_1)^{1/p} \tag{8}$$

is a metric on $\mathcal{M}_+(\mathbb{T})$.

Consider a situation where we need to discriminate between two signals on the basis of their statistics or of their power spectra. In such cases, additive or multiplicative noise typically impedes our ability to differentiate between the two. In particular, we have that $\Phi \mapsto \Phi + \Phi_{\mathrm{a}}$ represents the operation of adding independent noise with spectrum Φ_{a} and $\Phi \mapsto \Phi * \Phi_{\mathrm{m}}$ represents the operation of multiplying the signal with independent noise with spectrum Φ_{m}. This was considered in [27] and it was shown that the transportation distance respects this property in the sense that corrupting two signals with additive and (normalized) multiplicative noise decreases their transportation distance. Specifically, the following theorem holds.

Theorem 2 ([27])**.** *Let* $p \geq 1$, $\varkappa > 0$, *and let the cost function be* $c(\vartheta, \varphi) = |\vartheta - \varphi|^p$, *and let* $W_\varkappa(\Phi_0, \Phi_1)$ *be defined by* (8). *Then,* $W_\varkappa(\Phi_0, \Phi_1)$ *is contractive with respect to the additive and normalized multiplicative noise, i.e.,*

- $W_\varkappa(\Phi_0 + \Phi_{\mathrm{a}}, \Phi_1 + \Phi_{\mathrm{a}}) \leq W_\varkappa(\Phi_0, \Phi_1)$
- $W_\varkappa(\Phi_0 * \Phi_{\mathrm{m}}, \Phi_1 * \Phi_{\mathrm{m}}) \leq W_\varkappa(\Phi_0, \Phi_1)$,

for any $\Phi_{\mathrm{a}}, \Phi_{\mathrm{m}} \in \mathcal{M}_+(\mathbb{T})$, *with* $\int_{\mathbb{T}} \Phi_{\mathrm{m}}(\vartheta) d\vartheta = 1$.

As we shall see, it is possible to construct notions of distance on the space of positive semi-definite Toeplitz matrices that have properties similar to those stated in Theorem 2.

3 A notion of distance for Toeplitz matrices

As noted above, any positive semi-definite Toeplitz matrix $\mathbf{R}$ has at least one spectral representation, i.e., there exists at least one spectrum Φ that is consistent with it. Thus, we define the distance, T, between two positive semi-definite Toeplitz matrices, $\mathbf{R}_0$ and $\mathbf{R}_1$, as the minimum transportation cost, as measured by S, between spectra consistent with the respective matrices, i.e.,

$$\begin{aligned} T(\mathbf{R}_0, \mathbf{R}_1) \triangleq \min_{\Phi_0, \Phi_1 \in \mathcal{M}_+(\mathbb{T})} \quad & S(\Phi_0, \Phi_1) \\ \text{subject to} \quad & \Gamma(\Phi_j) = \mathbf{R}_j,\ j = 0, 1. \end{aligned} \tag{9}$$

Considering the definition of S in (6), this can equivalently be formulated as the convex optimization problem

$$\begin{aligned} T(\mathbf{R}_0, \mathbf{R}_1) = \min_{M \in \mathcal{M}_+(\mathbb{T}^2)} \; & \int_{\mathbb{T}^2} c(\vartheta, \varphi) M(\vartheta, \varphi) d\vartheta d\varphi \\ \text{subject to } \; & \Gamma\left(\int_{\mathbb{T}} M(\vartheta, \varphi) d\varphi \right) = \mathbf{R}_0 \\ & \Gamma\left(\int_{\mathbb{T}} M(\vartheta, \varphi) d\vartheta \right) = \mathbf{R}_1. \end{aligned} \tag{10}$$

Note that the formulation in (10) is only defined for covariance matrices with the same diagonal, i.e., the same r_0, as defined in (1), or, equivalently, covariance matrices whose consistent spectra have the same mass. However, in order to allow for mass differences, $T(\mathbf{R}_0, \mathbf{R}_1)$ can be generalized in analog with (7) as

$$\begin{aligned} T_\varkappa(\mathbf{R}_0, \mathbf{R}_1) \triangleq \min_{\Phi_j \in \mathcal{M}_+(\mathbb{T})} \; & S_\varkappa(\Phi_0, \Phi_1) \\ \text{subject to } \; & \Gamma(\Phi_j) = \mathbf{R}_j \text{ for } j = 0, 1, \end{aligned} \tag{11}$$

or, equivalently,

$$\begin{aligned} T_\varkappa(\mathbf{R}_0, \mathbf{R}_1) = \min_{\substack{M \in \mathcal{M}_+(\mathbb{T}^2) \\ \Psi_0, \Psi_1 \in \mathcal{M}_+(\mathbb{T})}} \; & \int_{\mathbb{T}^2} c(\vartheta, \varphi) M(\vartheta, \varphi) d\vartheta d\varphi \\ & + \varkappa \left\| \int_{\mathbb{T}} M(\vartheta, \varphi) d\varphi - \Psi_0 \right\|_1 \\ & + \varkappa \left\| \int_{\mathbb{T}} M(\vartheta, \varphi) d\vartheta - \Psi_1 \right\|_1 \\ \text{subject to } \; & \Gamma(\Psi_j) = \mathbf{R}_j \text{ for } j = 0, 1. \end{aligned} \tag{12}$$

Typically, the cost function c would be selected to be symmetric in its arguments, in which case $T_\varkappa$ would be symmetric as well, i.e., $T_\varkappa(\mathbf{R}_0, \mathbf{R}_1) = T_\varkappa(\mathbf{R}_1, \mathbf{R}_0)$ for any Toeplitz matrices $\mathbf{R}_0, \mathbf{R}_1 \in \mathbb{M}_+^n$. Although many possible choices of such cost functions exist, we will in most examples presented herein consider $c(\vartheta, \varphi) = \left| e^{i\vartheta} - e^{i\varphi} \right|^2$, i.e., the cost function quantifies distances as the square of the distance between the corresponding points on the unit circle. As we show in Section 5, this particular choice of cost function allows for a sum-of-squares relaxation of the dual formulation of (12). This dual formulation is presented next.

3.1 The dual formulation

In order to study properties of the distance notion $T_\varkappa$, we consider the dual formulation of (12), where we assume that the cost function, c, is a continuous non-negative function on $\mathbb{T}^2$. In order to address this problem, we first note that the adjoint operator[2] $\Gamma^* : \mathbb{M}^n \to C_{\text{perio}}(\mathbb{T})$ of the operator Γ is

$$\Gamma^*(\mathbf{R})(\vartheta) = \frac{1}{2\pi}\mathbf{a}(\vartheta)^H \mathbf{R}\mathbf{a}(\vartheta)$$

since

$$\begin{aligned}\langle \Gamma(\Phi), \mathbf{R}\rangle &= \langle \frac{1}{2\pi}\int_{\mathbb{T}} \mathbf{a}(\vartheta)\Phi(\vartheta)\mathbf{a}(\vartheta)^H d\vartheta, \mathbf{R}\rangle \\ &= \int_{\mathbb{T}} \frac{1}{2\pi}\Phi(\vartheta)\mathbf{a}(\vartheta)^H \mathbf{R}\mathbf{a}(\vartheta) d\vartheta = \langle \Phi, \Gamma^*(\mathbf{R})\rangle,\end{aligned}$$

where, in the first line, the inner product is the one associated with $\mathbb{M}^n$, and the second line is the bilinear form with arguments $\Gamma^*(\mathbf{R}) \in C_{\text{perio}}(\mathbb{T})$ and $\Phi \in \mathcal{M}_+(\mathbb{T})$. With this result, we can derive an expression of the dual problem by considering the Lagrangian relaxation of (12). The Lagrangian is given by

$$\begin{aligned}\mathcal{L}_\varkappa(M, \Psi_0, \Psi_1, \mathbf{\Lambda}_0, \mathbf{\Lambda}_1) &= \int_{\mathbb{T}^2} c(\vartheta,\varphi)M(\vartheta,\varphi)d\vartheta d\varphi \\ &\quad + \langle \mathbf{R}_0 - \Gamma(\Psi_0), \mathbf{\Lambda}_0\rangle + \langle \mathbf{R}_1 - \Gamma(\Psi_1), \mathbf{\Lambda}_1\rangle \\ &\quad + \varkappa\left\|\int_{\mathbb{T}} M(\vartheta,\varphi)d\varphi - \Psi_0\right\|_1 + \varkappa\left\|\int_{\mathbb{T}} M(\vartheta,\varphi)d\vartheta - \Psi_1\right\|_1 \\ &= \langle \mathbf{\Lambda}_0, \mathbf{R}_0\rangle + \langle \mathbf{\Lambda}_1, \mathbf{R}_1\rangle + \int_{\mathbb{T}^2} c(\vartheta,\varphi)M(\vartheta,\varphi)d\vartheta d\varphi \\ &\quad - \langle \Psi_0, \Gamma^*(\Lambda_0)\rangle - \langle \Psi_1, \Gamma^*(\Lambda_1)\rangle \\ &\quad + \varkappa\left\|\int_{\mathbb{T}} M(\vartheta,\varphi)d\varphi - \Psi_0\right\|_1 + \varkappa\left\|\int_{\mathbb{T}} M(\vartheta,\varphi)d\vartheta - \Psi_1\right\|_1.\end{aligned}$$

Note that the Lagrange multiplier matrices $\mathbf{\Lambda}_0$ and $\mathbf{\Lambda}_1$ may be taken as Hermitian matrices, as $\mathbf{R}_0 - \Gamma(\Psi_0)$ and $\mathbf{R}_1 - \Gamma(\Psi_1)$ are Hermitian, and thus all inner products are real. Considering the infimum of $\mathcal{L}_\varkappa$ with respect to Ψ_0 and Ψ_1, it may be noted that this is only finite if $\Gamma^*(\Lambda_0)(\vartheta) \leq \varkappa$ and $\Gamma^*(\Lambda_1)(\varphi) \leq \varkappa$, for

[2] Strictly speaking, this is the pre-adjoint operator of Γ.

all $\vartheta, \varphi \in \mathbb{T}$. If this is satisfied, the Lagrangian is, for any fixed non-negative M, minimized by Ψ_0 and Ψ_1 given by

$$\Psi_0(\vartheta) = \mathbb{1}_{\{\Gamma^*(\Lambda_0)(\vartheta) \in [-\varkappa, \varkappa]\}} \int_{\mathbb{T}} M(\vartheta, \varphi) d\varphi$$
$$\Psi_1(\varphi) = \mathbb{1}_{\{\Gamma^*(\Lambda_1)(\varphi) \in [-\varkappa, \varkappa]\}} \int_{\mathbb{T}} M(\vartheta, \varphi) d\vartheta,$$

where $\mathbb{1}_{\{\cdot\}}$ is the indicator function. Using this, and considering the infimum with respect to M, we arrive at

$$\inf_{\substack{M \in \mathcal{M}_+(\mathbb{T}^2) \\ \Psi_0, \Psi_1 \in \mathcal{M}_+(\mathbb{T})}} \mathcal{L}_\varkappa = \begin{cases} \langle \mathbf{\Lambda}_0, \mathbf{R}_0 \rangle + \langle \mathbf{\Lambda}_1, \mathbf{R}_1 \rangle & \text{if } (\mathbf{\Lambda}_0, \mathbf{\Lambda}_1) \in \Omega_c^\varkappa \\ -\infty & \text{otherwise} \end{cases}$$

where

$$\begin{aligned} \Omega_c^\varkappa = \{ & \mathbf{\Lambda}_0, \mathbf{\Lambda}_1 \in \mathbb{M}^n \mid \\ & \Gamma^*(\mathbf{\Lambda}_0)(\vartheta) + \Gamma^*(\mathbf{\Lambda}_1)(\varphi) \leq c(\vartheta, \varphi), \\ & \Gamma^*(\mathbf{\Lambda}_0)(\vartheta) \leq \varkappa, \ \Gamma^*(\mathbf{\Lambda}_1)(\varphi) \leq \varkappa \ \text{ for all } \vartheta, \varphi \in \mathbb{T} \}, \end{aligned}$$

where we have used that the cost function c is non-negative. This yields the dual problem.

Proposition 1. *Let the cost function c be continuous and non-negative, and let $\varkappa > 0$. Then, the dual problem of* (12) *is*

$$\underset{(\mathbf{\Lambda}_0, \mathbf{\Lambda}_1) \in \Omega_c^\varkappa}{\text{maximize}} \quad \langle \mathbf{\Lambda}_0, \mathbf{R}_0 \rangle + \langle \mathbf{\Lambda}_1, \mathbf{R}_1 \rangle, \tag{13}$$

where

$$\begin{aligned} \Omega_c^\varkappa = \{ & \mathbf{\Lambda}_0, \mathbf{\Lambda}_1 \in \mathbb{M}^n \mid \\ & \Gamma^*(\mathbf{\Lambda}_0)(\vartheta) + \Gamma^*(\mathbf{\Lambda}_1)(\varphi) \leq c(\vartheta, \varphi), \\ & \Gamma^*(\mathbf{\Lambda}_0)(\vartheta) \leq \varkappa, \ \Gamma^*(\mathbf{\Lambda}_1)(\varphi) \leq \varkappa \ \ \forall \vartheta, \varphi \in \mathbb{T} \}. \end{aligned}$$

Since the primal and dual problems are convex and the set of feasible points $\Omega_c^\varkappa$ has non-empty interior for any $\varkappa > 0$, Slater's condition (see, e.g., [41]) gives that strong duality holds and hence the duality gap between (12) and (13) is

zero. This can be generalized to hold for lower semi-continuous cost functions analogous to [42, proof of Theorem 4.1]. Also, note the strong resemblance in form compared to the dual of (7), which is given by [27]

$$\begin{aligned}
\underset{\lambda_0,\lambda_1\in C_{\text{perio}}(\mathbb{T})}{\text{maximize}} \quad & \int_{\mathbb{T}} \lambda_0(\vartheta)\Phi_0(\vartheta)d\vartheta + \int_{\mathbb{T}} \lambda_1(\varphi)\Phi_1(\varphi)d\varphi \\
\text{subject to} \quad & \lambda_0(\vartheta)+\lambda_1(\varphi) \leq c(\vartheta,\varphi) \quad \text{for all } \vartheta,\varphi\in\mathbb{T} \\
& \lambda_0(\vartheta) \leq \varkappa \quad \text{for all } \vartheta\in\mathbb{T} \\
& \lambda_1(\varphi) \leq \varkappa \quad \text{for all } \varphi\in\mathbb{T}.
\end{aligned}$$

Similarly, for the case when the diagonals of $\mathbf{R}_0$ and $\mathbf{R}_1$ are required to be equal, the dual problem is

$$\underset{(\mathbf{\Lambda}_0,\mathbf{\Lambda}_1)\in\Omega_c}{\text{maximize}} \quad \langle \mathbf{\Lambda}_0, \mathbf{R}_0\rangle + \langle \mathbf{\Lambda}_1, \mathbf{R}_1\rangle, \tag{14}$$

where

$$\Omega_c = \{\mathbf{\Lambda}_0, \mathbf{\Lambda}_1 \in \mathbb{M}^n \mid \Gamma^*(\mathbf{\Lambda}_0)(\vartheta) + \Gamma^*(\mathbf{\Lambda}_1)(\varphi) \leq c(\vartheta,\varphi) \ \ \forall \vartheta,\varphi\in\mathbb{T}\}.$$

This is the dual problem of (9), where adding and subtracting mass is not allowed in the transport problem.

3.2 Properties of distance notion $T_\varkappa$

For the distance notion $T_\varkappa$ in (11), the following proposition holds.

Proposition 2. *Let $\varkappa > 0$ and let the cost function c be a continuous function and a semi-metric on $\mathbb{T}$. Then, the distance notion $T_\varkappa$ in* (11) *is a semi-metric on the set of positive semi-definite Toeplitz matrices.*

Proof. See appendix. □

The implication of the semi-metric property is that $T_\varkappa$ may indeed be used to quantify distances between covariance matrices, or, stochastic processes. We may also state the following proposition.

Proposition 3. *Let the cost function c be continuous and non-negative and such that $c(\vartheta,\vartheta) = 0$, for all $\vartheta \in \mathbb{T}$. Then, the distance measure $T_\varkappa$ defined in* (11) *is contractive with respect to additive noise. If c is also shift-invariant, i.e.,*

$c(\vartheta - \phi, \varphi - \phi) = c(\vartheta, \varphi)$, for all $\vartheta, \varphi, \phi \in \mathbb{T}$, then $T_\varkappa$ is also contractive with respect to multiplicative noise whose covariance matrix has diagonal elements smaller than or equal to unity.

Proof. See appendix. □

The statement of this proposition is that when two stochastic processes become contaminated by noise, the distance notion $T_\varkappa$ decreases and hence the processes become harder to distinguish. Intuitively, this is a desirable property of $T_\varkappa$, as additive or multiplicative noise should indeed impede ones ability to discriminate between two processes. Proposition 3 may be proven by utilizing results in [27]. However, in the interest of making the exposition self-contained, we provide a direct proof in the appendix based on the dual formulation in (13). It may be noted that the assumptions made in Proposition 3 regarding the cost function c are quite mild, allowing for a large class of potential cost functions. In particular, $c(\vartheta, \varphi) = \left|e^{i\vartheta} - e^{i\varphi}\right|^2$ satisfies the conditions of Proposition 3 for both additive and multiplicative noise.

4 Interpolation, extrapolation, and tracking

The formulation in (9) does not only define a notion of distance between two Toeplitz covariance matrices, $\mathbf{R}_0$ and $\mathbf{R}_1$; it also provides a means of forming interpolating matrices, i.e., defining intermediate covariance matrices $\mathbf{R}_\tau$, for $\tau \in (0, 1)$. In order to define the interpolating matrices, we again utilize the spectral representation of positive semi-definite Toeplitz matrices. To this end, we note that, given an optimal transport plan, M, found as the functional minimizing (9), one may define spectra intermediate to the marginals $\int_{\mathbb{T}} M(\vartheta, \varphi) d\varphi$ and $\int_{\mathbb{T}} M(\vartheta, \varphi) d\vartheta$ by linearly shifting the frequency locations of the spectral mass, as dictated by M. That is, any mass transferred from ϕ to $\phi + \varphi$ is defined to, at $\tau \in [0, 1]$, be located at frequency $\phi + \tau\varphi$. Using this, the intermediate spectrum is given by

$$\begin{aligned} \Phi_\tau^M(\vartheta) &\triangleq \int_{\mathbb{T}^2} \delta_\vartheta(\{\phi + \tau\varphi\}_{\mathrm{mod}\,\mathbb{T}}) M(\phi, \phi + \varphi) d\phi d\varphi \\ &= \int_{\mathbb{T}} M(\vartheta - \tau\varphi, \vartheta + (1 - \tau)\varphi) d\varphi. \end{aligned} \tag{15}$$

Here, δ_ϑ is the Dirac delta function localized at ϑ, i.e., $\delta_\vartheta(\phi) \triangleq \delta(\phi - \vartheta)$ and the integrands are extended periodically with period 2π outside the domain of integration. Also, we denote by $\{x\}_{\mathrm{mod}\,\mathbb{T}}$ the value in $\mathbb{T}$ that is congruent with x modulo 2π. Based on this definition of the interpolating spectrum Φ_τ^M, the corresponding interpolating covariance matrix, $\mathbf{R}_\tau$ is defined as the unique Toeplitz matrix consistent with this spectrum, i.e.,

$$\begin{aligned}\mathbf{R}_\tau &\triangleq \Gamma(\Phi_\tau^M) \\ &= \frac{1}{2\pi}\int_{\mathbb{T}} \mathbf{a}(\vartheta)\left(\int_{\mathbb{T}} M(\vartheta - \tau\varphi, \vartheta + (1-\tau)\varphi)d\varphi\right)\mathbf{a}(\vartheta)^H d\vartheta \\ &= \frac{1}{2\pi}\int_{\mathbb{T}^2} \mathbf{a}(\{\phi+\tau\varphi\}_{\mathrm{mod}\,\mathbb{T}})M(\phi,\phi+\varphi)\mathbf{a}(\{\phi+\tau\varphi\}_{\mathrm{mod}\,\mathbb{T}})^H d\phi d\varphi\end{aligned} \tag{16}$$

for $\tau \in [0,1]$. To simplify the following exposition, let $\mathcal{I}_\tau(M) \triangleq \Gamma(\Phi_\tau^M)$ denote the linear operator in (16) that maps a transport plan to an interpolating covariance matrix, i.e., $\mathbf{R}_\tau = \mathcal{I}_\tau(M)$. It may be noted from (16) that $\mathcal{I}_\tau(M)$ is well-defined also for $\tau \neq [0,1]$, i.e., the formulation allows also for extrapolation. The following proposition follows directly from the definition in (16).

Proposition 4. *For any $\tau \in \mathbb{R}$, the following basic properties hold for $\mathbf{R}_\tau = \mathcal{I}_\tau(M)$:*

a) If $\mathbf{R}_0$ and $\mathbf{R}_1$ have the same diagonal, then it is also the diagonal of $\mathbf{R}_\tau$.

b) The matrix $\mathbf{R}_\tau$ is a Toeplitz matrix.

c) The matrix $\mathbf{R}_\tau$ is positive semi-definite.

Due to these properties, the proposed method offers a way of interpolating the covariances of, e.g., slowly varying time series as the interpolant $\mathbf{R}_\tau$ allows for modeling linear changes in the spectrum of the process.

Remark 1. The interpolation approach generalizes trivially to the formulation in (12) between the covariances $\Gamma(\Phi_0^M)$ and $\Gamma(\Phi_1^M)$ under the assumption that Ψ_0 and Ψ_1 are perturbations of Φ_0^M and Φ_1^M, respectively. In order to define an interpolation and extrapolation procedure from $\mathbf{R}_0$ to $\mathbf{R}_1$ where there is a cost $\varkappa$ for adding and subtracting mass, one may, along with the interpolation path $\mathcal{I}_\tau(M)$, linearly add the part corresponding to the added and subtracted mass, i.e.,

$$\mathbf{R}_\tau = \mathcal{I}_\tau(M) + (1-\tau)\Gamma(\Psi_0 - \Phi_0^M) + \tau\Gamma(\Psi_1 - \Phi_1^M). \tag{17}$$

Note that in this scenario, positive semi-definiteness cannot be guaranteed for the extrapolation case.

4.1 Comparison with other methods

The properties in Proposition 4 distinguish the proposed interpolant $\mathbf{R}_\tau$ from other proposed matrix geodesics. As an example, consider the interpolant induced by the Euclidean metric, i.e., the distance between two covariance matrices $\mathbf{R}_0$ and $\mathbf{R}_1$ is defined as $\|\mathbf{R}_0 - \mathbf{R}_1\|_F$. This yields interpolants that are formed as convex combinations of $\mathbf{R}_0$ and $\mathbf{R}_1$, i.e., $\mathbf{R}_\tau^{\text{conv}} = (1-\tau)\mathbf{R}_0 + \tau\mathbf{R}_1$, for $\tau \in [0,1]$. This preserves the Toeplitz structure, as well as the diagonal of the end-point matrices and the positive semi-definiteness. However, from a spectral representation point of view, the convex combination gives rise to fade-in fade-out effects, i.e., only spectral modes directly related to $\mathbf{R}_0$ and $\mathbf{R}_1$ can be represented, and there can be no shift in the location of these modes (see also Example 6.2 and Figure 3). Other more sophisticated options include, e.g., the geodesic with respect to g-convexity [9]

$$\tilde{\mathbf{R}}_\tau = \mathbf{R}_0^{1/2}\left(\mathbf{R}_0^{-1/2}\mathbf{R}_1\mathbf{R}_0^{-1/2}\right)^\tau \mathbf{R}_0^{1/2} \tag{18}$$

and the geodesic in [20, 21], which builds on optimal mass transport of Gaussian distributions and can be expressed as

$$\breve{\mathbf{R}}_\tau = \left((1-\tau)\mathbf{R}_0^{1/2} + \tau\mathbf{R}_1^{1/2}\mathbf{U}\right)\left((1-\tau)\mathbf{R}_0^{1/2} + \tau\mathbf{R}_1^{1/2}\mathbf{U}\right)^H, \tag{19}$$

where

$$\mathbf{U} = \mathbf{R}_1^{-1/2}\mathbf{R}_0^{-1/2}\left(\mathbf{R}_0^{1/2}\mathbf{R}_1\mathbf{R}_0^{1/2}\right)^{1/2}.$$

One may also perform interpolation using geodesics induced by the log-Euclidean metric (see, e.g., [43]), i.e., where distances are defined as $\|\log(\mathbf{R}_0) - \log(\mathbf{R}_1)\|_F$, with $\log(\cdot)$ here denoting the matrix logarithm. For this case, the geodesic is given by

$$\mathbf{R}_\tau^{\text{log-Euclid}} = \exp\left((1-\tau)\log(\mathbf{R}_0) + \tau\log(\mathbf{R}_1)\right), \tag{20}$$

where $\exp(\cdot)$ denotes the matrix exponential. It may here be noted, that although the three geodesics in (18), (19), and (20) preserve positive definiteness, they are not defined for singular matrices due to the use of matrix inverses and matrix logarithms. Also, for general Toeplitz covariance matrices, these geodesics preserve

neither the Toeplitz structure nor the diagonal of the end-point matrices. Further, as noted above, the three properties in Proposition 4 hold for any $\tau \in \mathbb{R}$ for the proposed approach, and thus directly allows for extrapolating using (16). In contrast, it may be noted that for the linear combination $\mathbf{R}_\tau^{\text{conv}}$ there are no guarantees that the resulting matrix is positive semi-definite if $\tau \notin [0,1]$. Also, note that the alternative geodesics in (18), (19), and (20) do not naturally generalize to extrapolation.

4.2 Tracking of slowly varying processes

The proposed interpolant $\mathbf{R}_\tau = \mathcal{I}_\tau(M)$ may also be readily used for tracking slowly varying stochastic process. As noted above, $\mathcal{I}_\tau(M)$ allows for the modeling of slow, i.e., locally linear, shifts in the location of spectral power. Building on this property, we can extend the optimal transport problem in (9), in order to fit a covariance path $\mathbf{R}_\tau$ to a sequence of J covariance matrix estimates, $\hat{\mathbf{R}}_{\tau_j}$, for $j = 1, \ldots, J$. As $\mathbf{R}_\tau$ is unambiguously determined from a transport plan M via $\mathcal{I}_\tau(M)$, this tracking problem may be formulated as the convex optimization problem

$$\underset{M \in \mathcal{M}_+(\mathbb{T}^2)}{\text{minimize}} \int_{\mathbb{T}^2} c(\vartheta,\varphi) M(\vartheta,\varphi) d\vartheta d\varphi + \lambda \sum_{j=1}^{J} \left\| \mathcal{I}_{\tau_j}(M) - \hat{\mathbf{R}}_{\tau_j} \right\|_F^2, \tag{21}$$

where $\lambda > 0$ is a user-specified regularization parameter. As may be noted from (21), the optimal transport plan M is here determined as the one that minimizes not only the transport cost, but also takes into account the deviations of the interpolant $\mathbf{R}_{\tau_j} = \mathcal{I}_{\tau_j}(M)$ from the available covariance matrix estimates $\hat{\mathbf{R}}_{\tau_j}$. The behavior of this construction is illustrated in the numerical section.

4.3 Clustering: Barycenter computation

As a further example, we will see that the barycenter with respect to the distance notion $T_\varkappa$ may be formulated as a convex optimization problem. This might be desirable in clustering or classification applications, where one is interested in either identifying classes of signals or processes based on their covariance matrices or associate a given covariance matrix with such a signal class. Considering the case of clustering, assume that L covariance matrices, $\mathbf{R}_\ell$, $\ell = 1, \ldots, L$, are avail-

able. Then, we may define their barycenter via $T_\varkappa$ according to

$$\mathbf{R}_{\text{bary}} = \underset{\mathbf{R}\in\mathbb{M}_+^n}{\arg\min} \sum_{\ell=1}^{L} T_\varkappa(\mathbf{R}, \mathbf{R}_\ell), \tag{22}$$

i.e., as the covariance matrix that minimizes the sum of $T_\varkappa$ for the set of covariance matrices $\mathbf{R}_\ell$. Explicitly, $\mathbf{R}_{\text{bary}}$ solves the convex optimization problem

$$\begin{aligned}
\underset{\substack{\mathbf{R}\in\mathbb{M}_+^n \\ M_\ell\in\mathcal{M}_+(\mathbb{T}^2) \\ \Psi_\ell\in\mathcal{M}_+(\mathbb{T}) \\ \Phi_\ell\in\mathcal{M}_+(\mathbb{T})}}{\text{minimize}} \quad & \sum_{\ell=1}^{L} \int_{\mathbb{T}^2} c(\vartheta,\varphi) M_\ell(\vartheta,\varphi) d\vartheta d\varphi \\
& + \varkappa \sum_{\ell=1}^{L} \left\| \int_{\mathbb{T}} M_\ell(\vartheta,\phi) d\vartheta - \Phi_\ell \right\|_1 \\
& + \varkappa \sum_{\ell=1}^{L} \left\| \int_{\mathbb{T}} M_\ell(\vartheta,\phi) d\phi - \Psi_\ell \right\|_1 \\
\text{subject to} \quad & \Gamma(\Phi_\ell) = \mathbf{R} \\
& \Gamma(\Psi_\ell) = \mathbf{R}_\ell, \ \ell = 1,\dots,L.
\end{aligned} \tag{23}$$

This formulation allows for using, e.g., K-means clustering (see, e.g., [44]) in order to identify classes of covariance matrices, as well as classify a given covariance matrix according to these classes. Classification of a covariance matrix $\mathbf{R}$ according to classes defined by a set of barycenters $\mathbf{R}_{\text{bary}}^{(j)}$, $j = 1,\dots,J$, may then be formulated as

$$\underset{j\in\{1,\dots,J\}}{\arg\min}\ T_\varkappa\left(\mathbf{R}, \mathbf{R}_{\text{bary}}^{(j)}\right). \tag{24}$$

In Section 6.4, we present a simple illustration of this potential application, considering unsupervised clustering of phonemes.

5 Sum-of-squares relaxation

In order to solve the dual problem in (13) in practice, it has to be implemented as a finite dimensional problem, e.g., by gridding the space $\mathbb{T}^2$ and thereby approximating the set $\Omega_c^\varkappa$ using a finite number of constraints. However, for the special

case of $c(\vartheta, \varphi) = \left|e^{i\vartheta} - e^{i\varphi}\right|^{2p}$, for $p \in \mathbb{N}$, this can also be done using a sum-of-squares (SOS) relaxation (see, e.g., [45]) of (13). For simplicity of notation, we present here the case with $c(\vartheta, \varphi) = \left|e^{i\vartheta} - e^{i\varphi}\right|^2$ and without $\varkappa$ ($\varkappa = \infty$), i.e., the dual of (9) as formulated in (14). To this end, identify

$$z = e^{i\vartheta},\ w = e^{i\varphi},$$

which allows us to write

$$\begin{aligned} c(\vartheta, \varphi) &= 2 - zw^{-1} - z^{-1}w \\ \Gamma^*(\mathbf{\Lambda}_0)(\vartheta) &= \frac{1}{2\pi} \sum_{k=1}^{n} \sum_{\ell=1}^{n} [\mathbf{\Lambda}_0]_{k,\ell}\, z^{\ell-k} \\ \Gamma^*(\mathbf{\Lambda}_1)(\varphi) &= \frac{1}{2\pi} \sum_{k=1}^{n} \sum_{\ell=1}^{n} [\mathbf{\Lambda}_1]_{k,\ell}\, w^{\ell-k}. \end{aligned}$$

Thus, the set of constraints defining the feasible set Ω_c is given by $\Gamma^*(\mathbf{\Lambda}_0)(\vartheta) + \Gamma^*(\mathbf{\Lambda}_1)(\varphi) \leq c(\vartheta, \varphi)$, for all $\vartheta, \varphi \in \mathbb{T}$, or, equivalently,

$$2 - zw^{-1} - z^{-1}w - \frac{1}{2\pi} \sum_{k,\ell=1}^{n} [\mathbf{\Lambda}_0]_{k,\ell} z^{\ell-k} - \frac{1}{2\pi} \sum_{k,\ell=1}^{n} [\mathbf{\Lambda}_1]_{k,\ell} w^{\ell-k} \geq 0. \tag{25}$$

Note that in the two-dimensional trigonometric polynomial (25), the coefficient for $z^{-k_1} w^{-k_2}$ is equal to

$$\begin{cases} 2 - \frac{1}{2\pi} \operatorname{diag}(\mathbf{\Lambda}_0 + \mathbf{\Lambda}_1)^T \mathbf{1} & \text{for } k_1 = k_2 = 0 \\ -\frac{1}{2\pi} \operatorname{diag}(\mathbf{\Lambda}_0, k_1)^T \mathbf{1} & \text{for } k_1 \in \mathbb{Z}_{n\setminus 0} \text{ and } k_2 = 0 \\ -\frac{1}{2\pi} \operatorname{diag}(\mathbf{\Lambda}_1, k_2)^T \mathbf{1} & \text{for } k_1 = 0 \text{ and } k_2 \in \mathbb{Z}_{n\setminus 0} \\ -1 & \text{for } k_1 = 1 \text{ and } k_2 = -1 \\ -1 & \text{for } k_1 = -1 \text{ and } k_2 = 1 \\ 0 & \text{otherwise,} \end{cases}$$

where $\mathbb{Z}_{n\setminus 0} = \{k \in \mathbb{Z} \mid |k| < n, k \neq 0\}$. Here, $\operatorname{diag}(\mathbf{X}, k)$ denotes the column vector containing the elements on the kth super-diagonal of the matrix $\mathbf{X}$, if $k > 0$, and the elements on the kth sub-diagonal, if $k < 0$, with $\mathbf{1}$ denoting a column vector of ones of appropriate dimension.

In order to formulate a computationally feasible optimization problem, we remove the non-negativity constraint for the two-dimensional polynomial in (25)

and instead impose that it should have a sum-of-squares representation [45], [46]. In particular, we impose that the polynomial in (25) should be of the form

$$P(z,w) = \left(\mathbf{z}^{-1}\right)^T \mathbf{Q}\mathbf{z}, \tag{26}$$

where $\mathbf{Q} \in \mathbb{M}^{m^2}$ is positive semi-definite,

$$\begin{aligned}\mathbf{z} &= \begin{bmatrix}1 & w & \cdots & w^{m-1}\end{bmatrix}^T \otimes \begin{bmatrix}1 & z & \cdots & z^{m-1}\end{bmatrix}^T \\ \mathbf{z}^{-1} &= \begin{bmatrix}1 & w^{-1} & \cdots & w^{-m+1}\end{bmatrix}^T \otimes \begin{bmatrix}1 & z^{-1} & \cdots & z^{-m+1}\end{bmatrix}^T\end{aligned}$$

and $\otimes$ is the Kronecker product. Note that any polynomial on this form is non-negative by definition, and, furthermore, for any non-negative polynomal P^*, there is a sequence of polynomials, P_m, on the form (26) such that $\|P_m - P^*\|_\infty \to 0$ as $m \to \infty$.

Next, note that the coefficients of P are associated to the elements of $\mathbf{Q}$ in (26) according to

$$p_{k_1,k_2} = \operatorname{tr}\left(\mathbf{T}_{k_1,k_2}\mathbf{Q}\right), \quad \text{for } -m+1 \le k_1, k_2 \le m-1,$$

where $\mathbf{T}_{k_1,k_2} = \mathbf{T}_{k_2} \otimes \mathbf{T}_{k_1}$, and $\mathbf{T}_k$ is the matrix with ones on the kth diagonal and zeros elsewhere. Putting these facts together by requiring that the polynomial (25) can be written as a sum-of-square (26), we approximate (14) by the semi-definite program (SDP)

$$\begin{aligned}
\underset{\substack{\mathbf{\Lambda}_0,\mathbf{\Lambda}_1\in\mathbb{M}^n \\ \mathbf{Q}\in\mathbb{M}^{m^2}}}{\text{maximize}} \quad & \langle \mathbf{\Lambda}_0, \mathbf{R}_0\rangle + \langle \mathbf{\Lambda}_1, \mathbf{R}_1\rangle \qquad (27)\\
\text{subject to} \quad & \mathbf{Q} \succeq 0 \\
& \operatorname{tr}(\mathbf{Q}) = 2 - \frac{1}{2\pi}\operatorname{diag}(\mathbf{\Lambda}_0 + \mathbf{\Lambda}_1)^T \mathbf{1} \\
& \operatorname{tr}\left(\mathbf{T}_{-1,1}\mathbf{Q}\right) = -1 \\
& \operatorname{tr}\left(\mathbf{T}_{1,-1}\mathbf{Q}\right) = -1 \\
& \operatorname{tr}\left(\mathbf{T}_{k_1,0}\mathbf{Q}\right) = -\frac{1}{2\pi}\operatorname{diag}(\mathbf{\Lambda}_0, k_1)^T \mathbf{1}\ ,\ k_1 \in \mathbb{Z}_{n\setminus 0} \\
& \operatorname{tr}\left(\mathbf{T}_{0,k_2}\mathbf{Q}\right) = -\frac{1}{2\pi}\operatorname{diag}(\mathbf{\Lambda}_1, k_2)^T \mathbf{1}\ ,\ k_2 \in \mathbb{Z}_{n\setminus 0} \\
& \operatorname{tr}\left(\mathbf{T}_{k_1,k_2}\mathbf{Q}\right) = 0\ ,\ \text{if } k_1k_2 = 1 \text{ or } 1 < |k_1k_2| \le (m-1)^2.
\end{aligned}$$

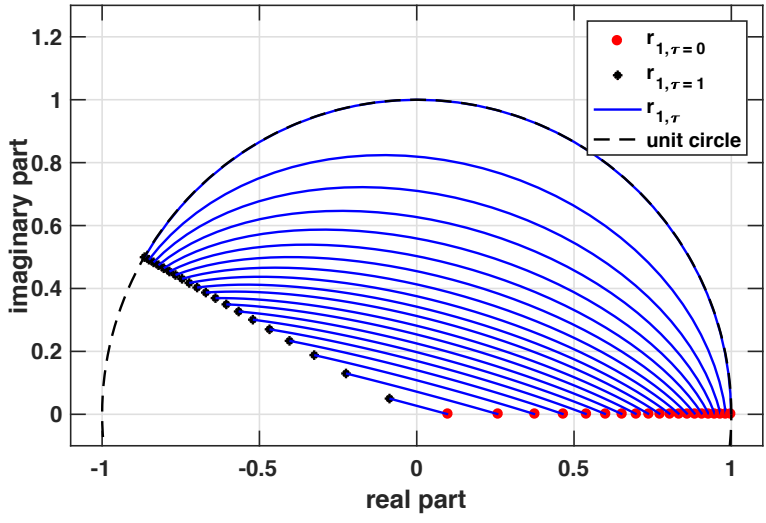

Figure 1: Path of the off-diagonal element of the covariance matrix $\mathbf{R}_\tau$ for the case $n = 2$, when choosing $\mathbf{R}_0$ and $\mathbf{R}_1$ according to (28).

It is worth noting that the P defined by the optimal $\mathbf{Q}$ in (26) is non-negative on the unit torus. Thus, for the solution of the SDP problem in (27), there will be a corresponding feasible solution of (14). Therefore, any solution to (27) will give a lower bound for the optimal objective value of (14). However, any non-negative polynomial may be arbitrarily well approximated by P of the form (26) by a suitable choice of the degree $m - 1$, and thus the maximal objective value of (27) will converge to the maximal objective value of (14) as m grows. By comparison, directly discretizing $\mathbb{T}$, thereby approximating Ω_c using a finite number of constraints, yields an upper bound for the maximal objective value of (14). In this case, the maximal objective value will also converge to that of (14), this time from above, as the spacing of the discretization of the grid becomes finer.

6 Numerical examples

In this section, we present some numerical examples illustrating different aspects and application areas of the proposed distance notion $T_\varkappa$, as well as the interpolant $\mathbf{R}_\tau = \mathcal{I}_\tau(M)$. Throughout these examples, we will use the cost function

$c(\vartheta, \varphi) = \left|e^{i\vartheta} - e^{i\varphi}\right|^2$. It should be stressed that many other choices of cost functions are possible, allowing for flexibility in modeling the specific scenario one is interested in. Apart from specific modeling aspects, one may preferably pick cost functions satisfying the assumptions in Propositions 2 and 3.

6.1 Trajectory example

In order to illustrate the behavior of the proposed interpolation method, we consider a simple scenario with covariance matrices of size $n = 2$. Consider covariance matrices of the form

$$\mathbf{R}_0 = \begin{bmatrix} 1 & r \\ r & 1 \end{bmatrix}, \ \mathbf{R}_1 = \begin{bmatrix} 1 & re^{i\frac{5\pi}{6}} \\ re^{-i\frac{5\pi}{6}} & 1 \end{bmatrix}, \tag{28}$$

where $r \in [-1, 1]$. Thus, considering interpolating paths $\mathbf{R}_\tau$, these will be on the form

$$\mathbf{R}_\tau = \begin{bmatrix} 1 & r_{1,\tau} \\ \bar{r}_{1,\tau} & 1 \end{bmatrix}. \tag{29}$$

Figure 1 displays the real and imaginary part of $r_{1,\tau}$ for $\tau \in [0, 1]$ when varying the magnitude r of the off-diagonal element of $\mathbf{R}_0$ and $\mathbf{R}_1$ between 0 and 1. As can be seen, the trajectories of $r_{1,\tau}$ approximately correspond to convex combinations $(1 - \tau)r + \tau re^{i\frac{5\pi}{6}}$ when r is close to zero, whereas they are considerably curved for r closer to 1. It may be noted that for the singular case, i.e., $r = 1$, the trajectory of $r_{1,\tau}$ coincides with the unit circle. Thus, we see that for covariance matrices that have consistent spectra that are essentially flat, the interpolant $\mathbf{R}_\tau$ will be approximately equal to the convex combination $(1 - \tau)\mathbf{R}_0 + \tau\mathbf{R}_1$. At the other extreme, the interpolant corresponding to covariance matrices $\mathbf{R}_0$ and $\mathbf{R}_1$ that have consistent spectra that are close to being singular will also have almost singular consistent spectra.

6.2 Interpolation and extrapolation for DOA

Next, we illustrate the proposed methods ability to produce interpolants $\mathbf{R}_\tau$ that are consistent with locally linear changes in the frequency location of spectral power. This is illustrated using a direction-of-arrival (DOA) estimation problem. Consider a uniform linear array (ULA) with 15 sensors with half-wavelength

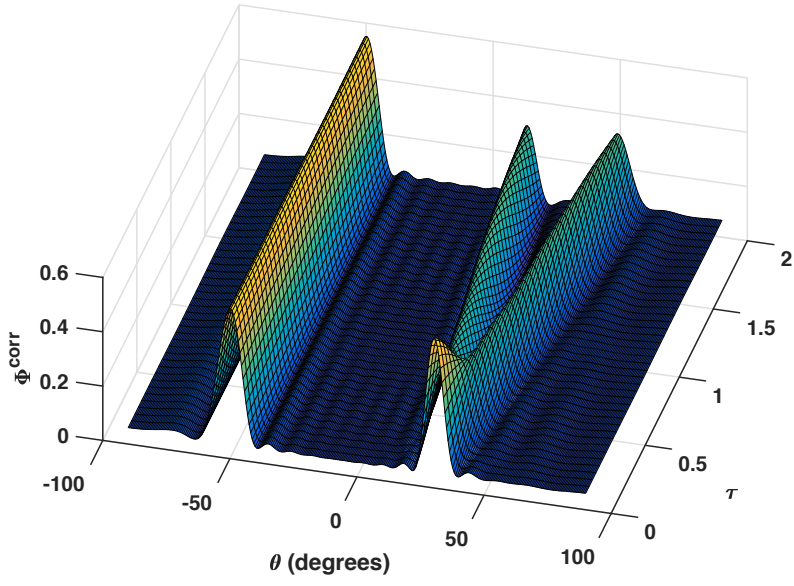

Figure 2: Interpolated spatial spectrum estimated as $\mathbf{a}(\vartheta)^H \mathbf{R}_\tau \mathbf{a}(\vartheta)$, where $\mathbf{R}_\tau$ is obtained by solving (10).

sensor spacing and a scenario where two covariance matrices

$$\mathbf{R}_0 = \frac{1}{2}\sum_{\ell=1}^{2} \mathbf{a}(\vartheta_\ell^{(0)})\mathbf{a}(\vartheta_\ell^{(0)})^H + \sigma^2\mathbf{I}$$

$$\mathbf{R}_1 = \frac{1}{2}\mathbf{a}(\vartheta_1^{(1)})\mathbf{a}(\vartheta_1^{(1)})^H + \frac{1}{4}\sum_{\ell=2}^{3} \mathbf{a}(\vartheta_\ell^{(1)})\mathbf{a}(\vartheta_\ell^{(1)})^H + \sigma^2\mathbf{I}$$

are available[3]. Here $\vartheta_1^{(0)} = \vartheta_1^{(1)} = -50°$, $\vartheta_2^{(0)} = 30°$, $\vartheta_2^{(1)} = 20°$, and $\vartheta_3^{(1)} = 40°$, and $\sigma^2 = 0.05$. Such a scenario may be interpreted as a target at $\vartheta_2^{(0)}$ splitting up into two targets at $\vartheta_2^{(1)}$ and $\vartheta_3^{(1)}$ as time progresses, whereas the target at $\vartheta_1^{(0)}$ stays put. We use the proposed method in (10), as $\mathbf{R}_0$ and $\mathbf{R}_1$ have the same diagonal, in order to find the optimal transport map M. Then, using

[3]Note that ϑ in this example denotes spatial frequency. For simplicity, we retain the notation $\mathbf{a}(\vartheta)$ also for this case.

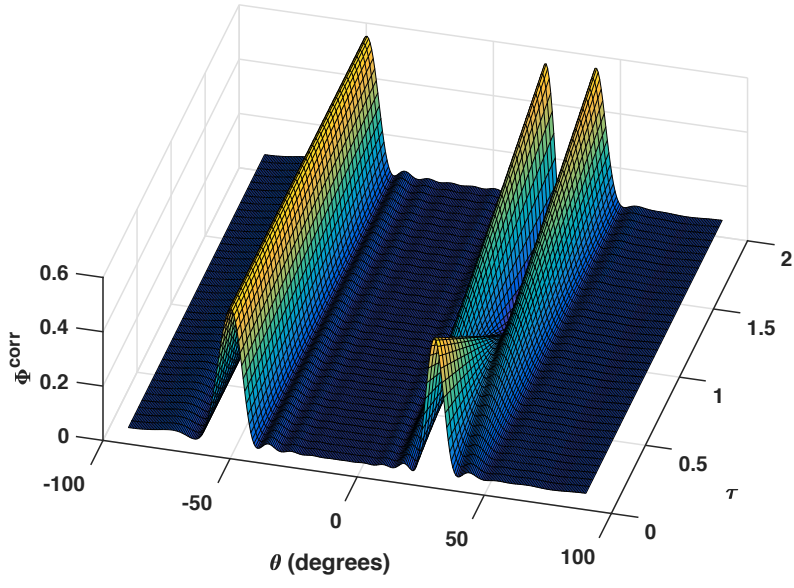

Figure 3: Interpolated spatial spectrum estimated as $\mathbf{a}(\vartheta)^H \mathbf{R}_\tau^{\text{conv}} \mathbf{a}(\vartheta)$, where $\mathbf{R}_\tau^{\text{conv}}$ is the linear combination of $\mathbf{R}_0$ and $\mathbf{R}_1$.

(16), we compute covariance matrices $\mathbf{R}_\tau$, for $\tau \in [0, 2]$, i.e., we both interpolate on $\tau \in [0, 1]$ and extrapolate on $\tau \in (1, 2]$. This is then compared to the basic interpolant $\mathbf{R}_\tau^{\text{conv}}$ based on convex and linear combinations of $\mathbf{R}_0$ and $\mathbf{R}_1$, as well as the more sophisticated covariance matrix geodesics $\tilde{\mathbf{R}}_\tau$ and $\breve{\mathbf{R}}_\tau$, as defined in (18) and (19), respectively. For these four cases, we then estimate the corresponding inter- and extrapolated spectra using the correlogram, i.e., as

$$\Phi^{\text{corr}}(\mathbf{X}, \vartheta) = \mathbf{a}(\vartheta)^H \mathbf{X} \mathbf{a}(\vartheta), \tag{30}$$

where $\mathbf{X}$ is substituted for the four different covariance interpolants. The results for the proposed interpolant $\mathbf{R}_\tau$ are shown in Figure 2. As can be seen, $\mathbf{R}_\tau$ indeed models a scenario where one of the targets has a constant location, whereas the second target splits upp into two smaller targets. Note also that $\mathbf{R}_\tau$ implies that the smaller target continue linearly with respect to the look-angle, ϑ, also for the extrapolation case, i.e., for $\tau > 1$. In contrast, the convex (linear for $\tau > 1$) combination $\mathbf{R}^{\text{conv}}$, as shown in Figure 3, displays undesirable behavior; clear

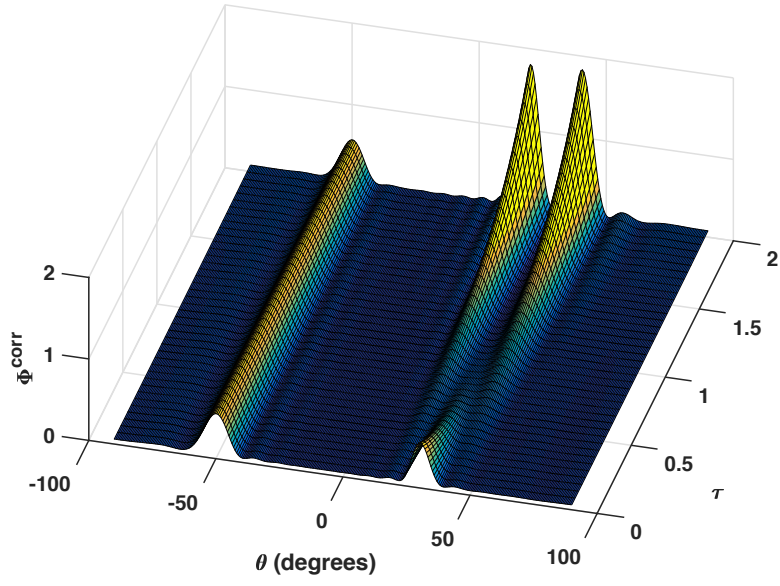

Figure 4: Interpolated spatial spectrum estimated as $\mathbf{a}(\vartheta)^H \tilde{\mathbf{R}}_\tau \mathbf{a}(\vartheta)$, where $\tilde{\mathbf{R}}_\tau$ is obtained from (18).

fade-in fade-out effects are visible, and non-negativity is violated as τ approaches 2 due to the fact that $\mathbf{R}^{\text{conv}}$ becomes indefinite. Similar objections may be raised against the geodesics $\tilde{\mathbf{R}}_\tau$ and $\breve{\mathbf{R}}_\tau$, shown in Figures 4 and 5, both displaying fade-in fade-out effects and thereby failing to model any displacement of the targets. Also, note that the total power of the signal varies greatly as τ goes from 0 to 2, especially for $\breve{\mathbf{R}}_\tau$.

6.3 Tracking of an AR-process

Next, we illustrate the approach in (21) for the tracking of signals with slowly varying spectral content. To this end, consider a complex autoregressive (AR) process with one complex, time-varying pole. The pole is placed at a constant radius of 0.9, and moves from the frequency 0.3π to 0.6π. Spectral estimates based directly on covariance matrix estimates $\hat{\mathbf{R}}$ are shown in the top plot of Figure 6. These covariance matrix estimates are obtained as the outer product estimate, based on 150 samples each, where the overlap between each estimate is

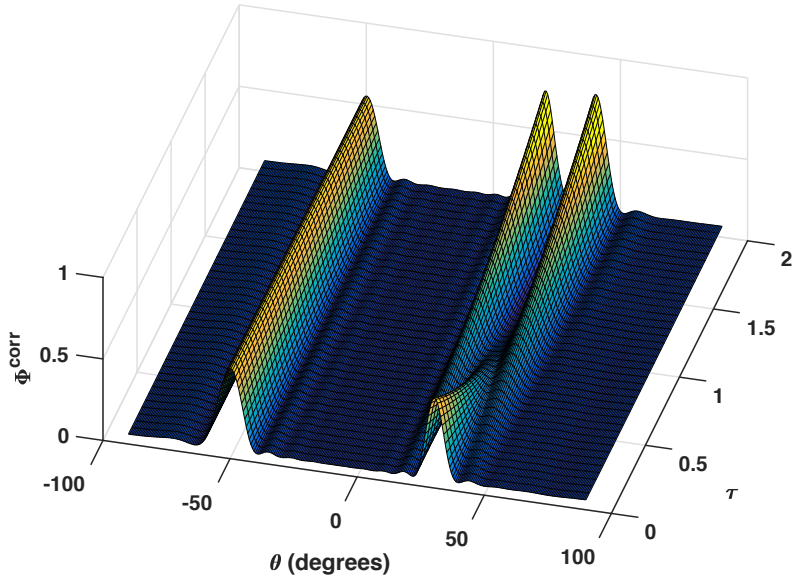

Figure 5: Interpolated spatial spectrum estimated as $\mathbf{a}(\vartheta)^H \breve{\mathbf{R}}_\tau \mathbf{a}(\vartheta)$, where $\breve{\mathbf{R}}_\tau$ is obtained from (19).

80 samples. Each estimated covariance matrix is of size $n = 15$. As can be seen, the spectral estimates are very noisy and vary greatly in power. Using five of these covariance matrix estimates $\hat{\mathbf{R}}$, evenly spaced throughout the signal, we solve (21) with $\lambda = \frac{1}{2n^2}$ in order to obtain an estimated covariance path, $\mathbf{R}_\tau$. The resulting spectra, estimated using (30), are shown in the bottom plot of Figure 6. As can be seen, the path resulting from the proposed method allows for a smooth tracking of the shift in spectral content. For comparison, Figure 7 displays the corresponding spectral estimates obtained from fitting geodesics induced by the Euclidean metric (in the top plot) as well as the log-Euclidean metric (in the bottom plot) to the same covariance estimates. Specifically, for the Euclidean case, the geodesic is constructed as $\mathbf{R}_\tau = (1-\tau)\mathbf{R}_0 + \tau\mathbf{R}_1$, where $(\mathbf{R}_0, \mathbf{R}_1)$ solves

$$\underset{\mathbf{R}_0 \succeq 0, \mathbf{R}_1 \succeq 0}{\arg\min} \sum_{j=1}^{J} \left\| (1-\tau_j)\mathbf{R}_0 + \tau_j \mathbf{R}_1 - \hat{\mathbf{R}}_{\tau_j} \right\|_F^2, \tag{31}$$

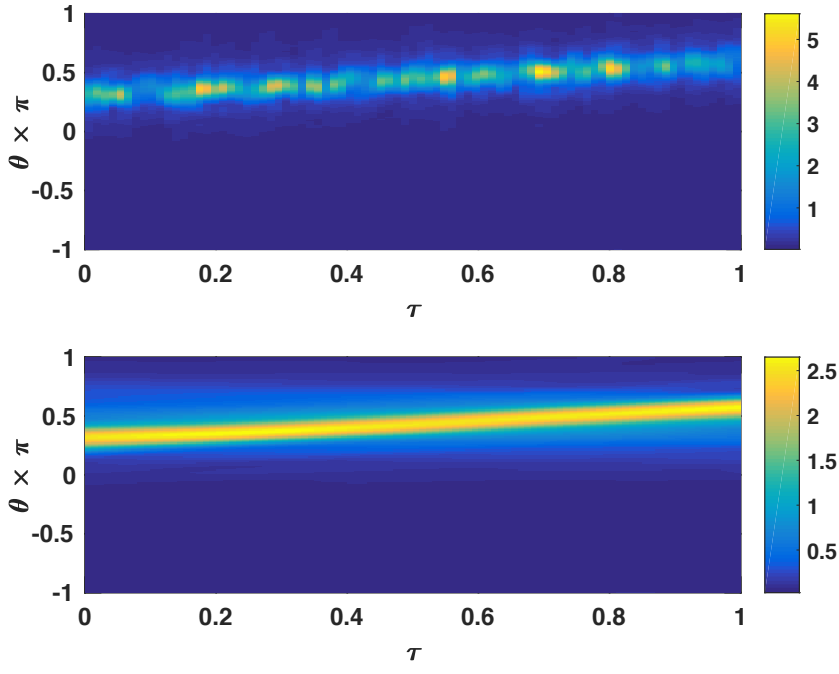

Figure 6: Spectrum estimated as $\mathbf{a}(\vartheta)^H \hat{\mathbf{R}} \mathbf{a}(\vartheta)$, where $\hat{\mathbf{R}}$ is estimated as the sample covariance matrix based on 150 samples in each window (top plot), as well as $\mathbf{a}(\vartheta)^H \mathbf{R}_\tau \mathbf{a}(\vartheta)$, where $\mathbf{R}_\tau$ is obtained by solving (21), fitted to a sequence of five covariance estimates (bottom plot).

whereas for the log-Euclidean metric (see, e.g., [43]), the geodesic is given by (20), where $(\mathbf{R}_0, \mathbf{R}_1)$ solves

$$\underset{\mathbf{R}_0 \succ 0, \mathbf{R}_1 \succ 0}{\arg\min} \sum_{j=1}^{J} \left\| (1-\tau_j) \log(\mathbf{R}_0) + \tau_j \log(\mathbf{R}_1) - \log\left(\hat{\mathbf{R}}_{\tau_j}\right) \right\|_F^2 . \tag{32}$$

As remarked earlier, we here require $\mathbf{R}_0$ and $\mathbf{R}_1$ to be positive definite in order for the geodesic to be defined. It may also be noted that this approach requires all estimated covariance matrices $\hat{\mathbf{R}}_{\tau_j}$ to have full rank. As may be seen in Figure 7, both these fitted geodesics imply spectral estimates displaying significant fade-in fade-out effects, which should be contrasted with the proposed method's ability to here produce a reasonable and intuitive interpolation.

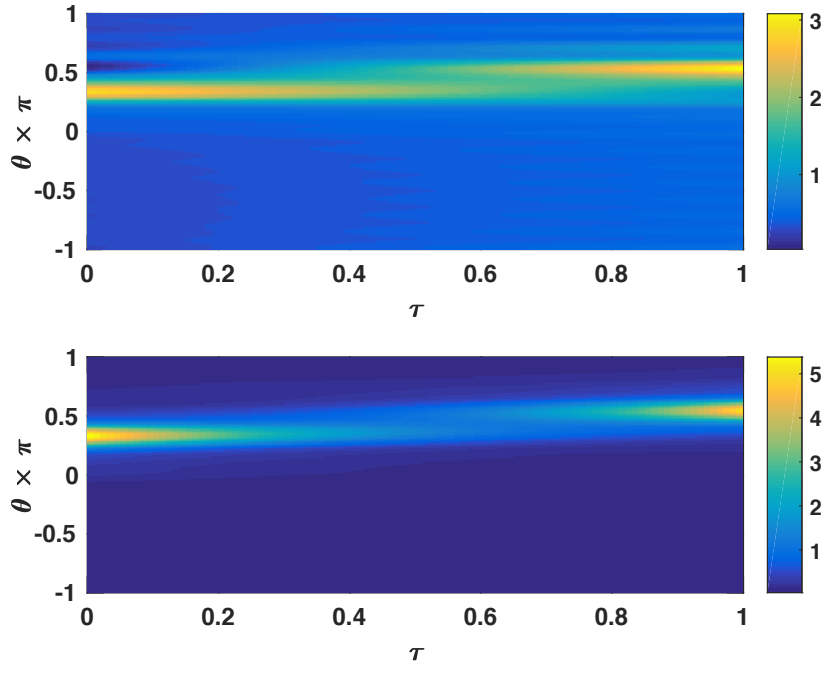

Figure 7: Spectrum estimated as $\mathbf{a}(\vartheta)^H \mathbf{R}_\tau \mathbf{a}(\vartheta)$, where $\mathbf{R}_\tau$ is the geodesic induced by the Euclidean metric (top plot) and the log-Euclidean metric (bottom plot), when fitted to a sequence of five covariance estimates.

6.4 K-means clustering

As a simple illustration of how to utilize the barycenter formulation in Section 4.3, we consider the application of unsupervised clustering of phonemes. Specifically, we consider 7 utterances; 3 utterances of the phoneme */ae/*, 2 utterances of */oy/*, and 2 utterances of */n/* taken from an annotated recording sampled at 16 kHz, with the durations of the different phonemes varying between 30 ms and 174 ms. For each utterance, we estimate an $n \times n$ covariance matrix of size $n = 10$, which is a quite common covariance matrix size in speech coding applications (see, e.g., [47]), and then run a K-means algorithm that alternates between classifying each estimated covariance matrix according to (24) and computing new barycenters according to (22). To ensure that the classification is unaffected by differences in signal power, which potentially could be a dominating factor, each covariance matrix estimate is normalized as to have its diagonal elements, i.e., r_0,

equal to one[4]. We initiate the algorithm by choosing initial barycenters as convex combinations of the available covariance matrix estimates, and demand separation into three clusters. The algorithm is then run until convergence, i.e., until the classification has stabilized. As a comparison, we perform the same K-means clustering using the Euclidean and the log-Euclidean metrics as described earlier, as well as the Kullback-Leibler divergence and the ellipticity distance measure introduced in [48] in order to compute distances as well as barycenters. Specifically, for two positive definite matrices $\mathbf{R}_0$ and $\mathbf{R}_1$ of size n, the Kullback-Leibler divergence is given by

$$d_{\mathrm{KL}}(\mathbf{R}_0, \mathbf{R}_1) = \mathrm{tr}\left(\mathbf{R}_0^{-1}\mathbf{R}_1\right) - \log\left|\mathbf{R}_0^{-1}\mathbf{R}_1\right| - n,$$

and the ellipticity distance is given by

$$d_{\mathrm{E}}(\mathbf{R}_0, \mathbf{R}_1) = n\log\left(\frac{1}{n}\mathrm{tr}\left(\mathbf{R}_0^{-1}\mathbf{R}_1\right)\right) - \log\left|\mathbf{R}_0^{-1}\mathbf{R}_1\right|.$$

As may be noted, these distance notions are only defined for non-singular matrices $\mathbf{R}_0$ and $\mathbf{R}_1$. For a set of J observed covariances $\mathbf{R}_j$, $j = 1, \ldots, J$, the barycenter induced by the Kullback-Leibler distance is given by

$$\mathbf{R}_{\mathrm{KL}}^{\mathrm{bary}} = \left(\frac{1}{J}\sum_{j=1}^{J}\mathbf{R}_j^{-1}\right)^{-1},$$

whereas for the ellipticity distance, the barycenter is the solution to the fixed point equation (see [48])

$$\mathbf{R}_{\mathrm{E}}^{\mathrm{bary}} = \left(\frac{n}{J}\sum_{j=1}^{J}\frac{\mathbf{R}_j^{-1}}{\mathrm{tr}\left(\mathbf{R}_j^{-1}\mathbf{R}_{\mathrm{E}}^{\mathrm{bary}}\right)}\right)^{-1}.$$

The solution $\mathbf{R}_{\mathrm{E}}^{\mathrm{bary}}$ is unique only up to a positive scaling factor, and thus, we here normalize $\mathbf{R}_{\mathrm{E}}^{\mathrm{bary}}$ as to have unit diagonal. It may be noted, that even for Toeplitz $\mathbf{R}_j$, $j = 1, \ldots, J$, the barycenters $\mathbf{R}_{\mathrm{KL}}^{\mathrm{bary}}$ and $\mathbf{R}_{\mathrm{E}}^{\mathrm{bary}}$ are not Toeplitz in general. As the K-means algorithm may converge to different solutions, i.e., dif-

[4]To preserve the diagonal elements of each matrix in this setting, any $\varkappa > \max_{\vartheta,\varphi} c(\vartheta, \varphi)$ constitues a valid choice.

Utterance	1	2	3	4	5	6	7
Phoneme	*/ae/*	*/ae/*	*/ae/*	*/oy/*	*/oy/*	*/n/*	*/n/*
Proposed	1	1	2	2	2	3	3
Euclidean	1	1	1	1	2	2	3
Log-Euclidean	1	1	1	1	2	2	3
KL divergence	1	1	2	1	3	3	3
Ellipticity	1	1	2	1	3	3	3

Table 1: Clustering of 7 utterances into three clusters using a K-means algorithm utilizing the barycenter formulation in Section 4.3, as well as four comparison distance measures. The third to seventh rows indicate the identified classes as given by the algorithm using the different distance measures.

Utterance	1	2	3	4	5	6	7
Class 1	1	1	4.20	7.10	5.02	10.54	18.67
Class 2	8.79	7.83	1	1	1	4.70	9.99
Class 3	20.51	24.93	8.30	12.05	1.40	1	1

Table 2: Distance, as measured by $T_\varkappa$, between each utterance and each barycenter for the three identified clusters. Each distance has been normalized by the least distance for each utterance.

ferent sets of clusters, depending on the choice of initial points, we have for each choice of distance measure run the algorithm several times using different starting points and selected the solution corresponding to the least total distance between each barycenters and its assigned covariance matrices. The results are shown in Table 1. As can be seen, the proposed distance notion $T_\varkappa$ produces a clustering in which the third instance of the */ae/* phoneme is erroneously grouped together with the two utterances of */oy/*; apart from this, the clustering corresponds well to the true phonemes. The comparison distance measures do in this example produce clusterings that differ more from the ground truth than the clustering using the proposed distance notion. It may also be noted that the Euclidean and log-Euclidean produce the same clustering. Table 2 presents the corresponding

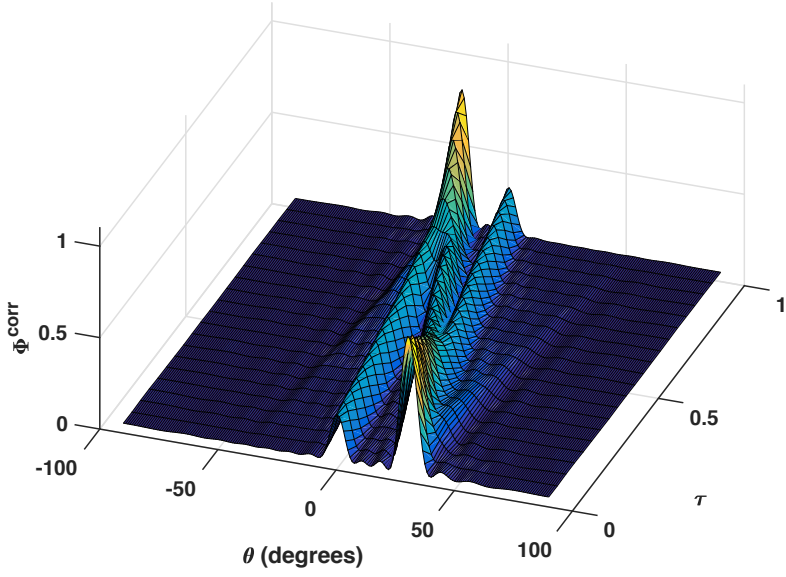

Figure 8: Interpolated spatial spectrum estimated as $\mathbf{a}(\vartheta)^H \mathbf{R}_\tau \mathbf{a}(\vartheta)$, where $\mathbf{R}_\tau$ is obtained by solving (10), for the case of one moving source and one static interferer, using with the cost function $c(\vartheta, \varphi) = \left| e^{i\vartheta} - e^{i\varphi} \right|^2$.

distance matrix, i.e., the matrix detailing the distance $T_\varkappa$ between each covariance matrix estimate and each barycenter. Note that for each utterance, the distances have been normalized by the least distance for that utterance. It is worth noting that the clusters are quite well-separated. Although being limited in scope, the example illustrates that the proposed distance notion may indeed be used in order to perform clustering and classification of stochastic processes based on their estimated covariances.

6.5 Fixed cost for moving mass

In some scenarios, there might be a relatively large noise component present in both $\mathbf{R}_0$ and $\mathbf{R}_1$, where the noise power is localized in frequency and this localization is the same in both $\mathbf{R}_0$ and $\mathbf{R}_1$. This might be the case in, e.g., DOA estimation scenarios where a source of interest is moving in the presence of a stationary interferer. Then, if the source for example moves past the location of the

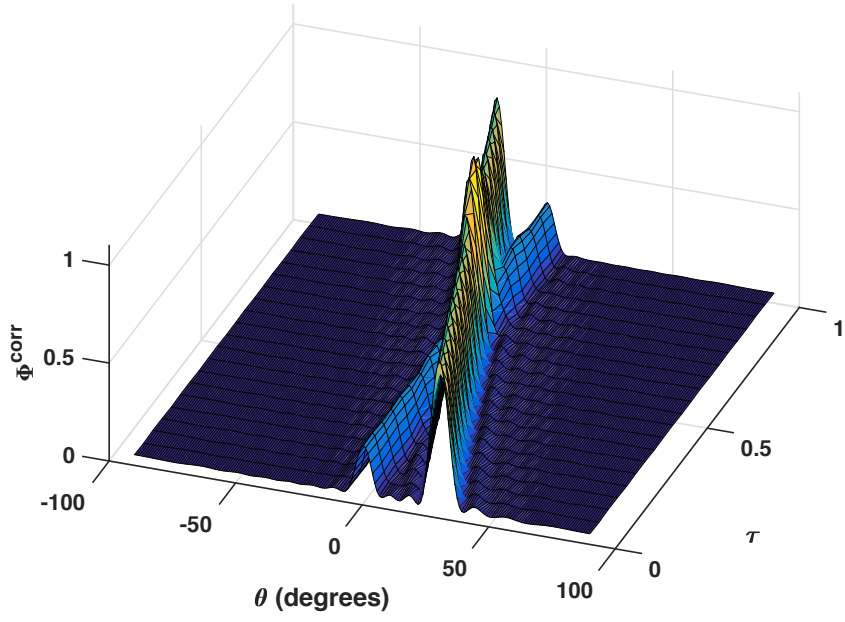

Figure 9: Interpolated spatial spectrum estimated as $\mathbf{a}(\vartheta)^H \mathbf{R}_\tau \mathbf{a}(\vartheta)$, where $\mathbf{R}_\tau$ is obtained by solving (10), for the case of one moving source and one static interferer, using with the modified cost function in (33).

interferer, the optimal transport problem in (10) may result in a power association such that the source and the interferer are mixed together. Such a scenario is illustrated in Figure 8, showing estimated spectra obtained from the interpolant $\mathbf{R}_\tau$, as given by (10) with the cost function $c(\vartheta, \varphi) = \left|e^{i\vartheta} - e^{i\varphi}\right|^2$. Here, the source is moving from look-angle $\vartheta = 30°$ to $\vartheta = -20°$, whereas there is a fixed interferer located at $\vartheta = 0°$ having a third of the power of the source. In order to avoid this type of problem and promoting transport plans that avoid transporting stationary masses, the cost function may be modified according to

$$c(\vartheta, \varphi) = \begin{cases} 1 + |e^{i\vartheta} - e^{i\varphi}|^2 & \text{if } \vartheta \neq \varphi \\ 0 & \text{if } \vartheta = \varphi. \end{cases} \tag{33}$$

Thus, the cost function is here formulated such that there is a fixed, baseline cost of moving any mass. The resulting estimated spectra obtained from the interpolant $\mathbf{R}_\tau$ resulting from solving (10) using this modified cost function is

shown in Figure 9. As can be seen, the source and the interferer are now well separated throughout the interpolated path.

7 Conclusions and further directions

In this work, we have proposed a notion of distance for positive semi-definite Toeplitz matrices. By considering spectral representations of such matrices, the proposed measure is based on distances, in an optimal mass transport sense, between families of spectra consistent with the Toeplitz matrices. We have shown that the proposed distance measure, under some mild assumptions, is contractive with respect to additive and multiplicative noise. The proposed measure may be used to, for example, define inter- and extrapolation of Toeplitz matrices, being of interest in applications such as tracking of slowly varying signals.

A future direction of this methodology is to generalize the distance measure for structured matrices such as Toeplitz-block-Toeplitz matrices and input-to-state covariances in the THREE framework for spectral analysis [12]. The latter generalizes ideas from beamspace processing, enabling the user to improve resolution and robustness in power spectral estimation over selected frequency bands [49, 50]. This will be the subject of further research.

8 Appendix

8.1 Proof of Proposition 2

Positivity and symmetry of $T_\varkappa$ follows directly from positivity and symmetry of c. Further, is clear that $T_\varkappa(\mathbf{R}_0, \mathbf{R}_1) = 0$ if $\mathbf{R}_0 = \mathbf{R}_1$. Thus, it remains to show that $T_\varkappa(\mathbf{R}_0, \mathbf{R}_1) = 0$ implies that $\mathbf{R}_0 = \mathbf{R}_1$. Next, if the objective function (12) is equal to zero, then the first term is zero and since c is a semi-metric, the transport plan M only has support on $\vartheta = \varphi$. Since the second and third terms of (12) are zero, it follows that $\Psi_0 = \Psi_1$, and hence $\mathbf{R}_0 = \mathbf{R}_1$. □

8.2 Proof of Proposition 3

First, we show that $T_\varkappa$ is contractive with respect to additive noise. Consider two processes with covariance matrices $\mathbf{R}_0$ and $\mathbf{R}_1$ and assume that they are both additively corrupted by an independent noise process with covariance $\mathbf{R}_w$. This

results in processes with covariances $\mathbf{R}_0' = \mathbf{R}_0 + \mathbf{R}_w$ and $\mathbf{R}_1' = \mathbf{R}_1 + \mathbf{R}_w$, respectively. From the dual formulation in (13), the distance between these covariance matrices is given by

$$T_\varkappa(\mathbf{R}_0', \mathbf{R}_1') = \max_{(\mathbf{\Lambda}_0, \mathbf{\Lambda}_1) \in \Omega_c^\varkappa} \langle \mathbf{\Lambda}_0, \mathbf{R}_0 \rangle + \langle \mathbf{\Lambda}_1, \mathbf{R}_1 \rangle + \langle \mathbf{\Lambda}_0 + \mathbf{\Lambda}_1, \mathbf{R}_w \rangle.$$

Let Φ_w be any spectrum consistent with the noise covariance $\mathbf{R}_w$, i.e., with $\Gamma(\Phi_w) = \mathbf{R}_w$. We then have that

$$\begin{aligned} \langle \mathbf{\Lambda}_0 + \mathbf{\Lambda}_1, \mathbf{R}_w \rangle &= \langle \mathbf{\Lambda}_0 + \mathbf{\Lambda}_1, \Gamma(\Phi_w) \rangle \\ &= \langle \Gamma^*(\mathbf{\Lambda}_0) + \Gamma^*(\mathbf{\Lambda}_1), \Phi_w \rangle \\ &= \int_{\mathbb{T}} \left(\Gamma^*(\mathbf{\Lambda}_0)(\vartheta) + \Gamma^*(\mathbf{\Lambda}_1)(\vartheta) \right) \Phi_w(\vartheta) d\vartheta. \end{aligned}$$

As $(\mathbf{\Lambda}_0, \mathbf{\Lambda}_1) \in \Omega_c^\varkappa$, it holds that

$$\Gamma^*(\mathbf{\Lambda}_0)(\vartheta) + \Gamma^*(\mathbf{\Lambda}_1)(\vartheta) \leq c(\vartheta, \vartheta) = 0 \text{ for all } \vartheta \in \mathbb{T}$$

and since $\Phi_w \geq 0$, we get

$$\langle \mathbf{\Lambda}_0 + \mathbf{\Lambda}_1, \mathbf{R}_w \rangle \leq 0.$$

Hence, it follows that

$$T_\varkappa(\mathbf{R}_0', \mathbf{R}_1') \leq \max_{(\mathbf{\Lambda}_0, \mathbf{\Lambda}_1) \in \Omega_c} \langle \mathbf{\Lambda}_0, \mathbf{R}_0 \rangle + \langle \mathbf{\Lambda}_1, \mathbf{R}_1 \rangle = T_\varkappa(\mathbf{R}_0, \mathbf{R}_1).$$

Next, we show that $T_\varkappa$ is contractive with respect to multiplicative noise. Let the noise covariance matrix be $\mathbf{R}_w$, implying that the covariances of the contaminated processes are $\mathbf{R}_0' = \mathbf{R}_0 \odot \mathbf{R}_w$ and $\mathbf{R}_1' = \mathbf{R}_1 \odot \mathbf{R}_w$, respectively. Let Φ_w be any spectrum consistent with $\mathbf{R}_w$, i.e., $\Gamma(\Phi_w) = \mathbf{R}_w$. Also, let the diagonal elements of $\mathbf{R}_w$ be smaller than or equal to 1, so that

$$\int_{\mathbb{T}} \Phi_w(\vartheta) d\vartheta \leq 2\pi.$$

We have

$$\begin{aligned} T_\varkappa(\mathbf{R}_0', \mathbf{R}_1') &= \max_{(\mathbf{\Lambda}_0, \mathbf{\Lambda}_1) \in \Omega_c^\varkappa} \quad \langle \mathbf{\Lambda}_0, \mathbf{R}_0 \odot \mathbf{R}_w \rangle + \langle \mathbf{\Lambda}_1, \mathbf{R}_1 \odot \mathbf{R}_w \rangle \\ &= \max_{(\mathbf{\Lambda}_0, \mathbf{\Lambda}_1) \in \Omega_c^\varkappa} \quad \langle \mathbf{\Lambda}_0 \odot \overline{\mathbf{R}_w}, \mathbf{R}_0 \rangle + \langle \mathbf{\Lambda}_1 \odot \overline{\mathbf{R}_w}, \mathbf{R}_1 \rangle \\ &= \max_{\substack{(\mathbf{\Lambda}_0, \mathbf{\Lambda}_1) \in \Omega_c^\varkappa \\ \tilde{\mathbf{\Lambda}}_0, \tilde{\mathbf{\Lambda}}_1 \in \mathbb{M}^n}} \langle \tilde{\mathbf{\Lambda}}_0, \mathbf{R}_0 \rangle + \langle \tilde{\mathbf{\Lambda}}_1, \mathbf{R}_1 \rangle \\ &\qquad \text{subject to} \qquad \tilde{\mathbf{\Lambda}}_0 = \mathbf{\Lambda}_0 \odot \overline{\mathbf{R}_w} \\ &\qquad\qquad\qquad\quad \tilde{\mathbf{\Lambda}}_1 = \mathbf{\Lambda}_1 \odot \overline{\mathbf{R}_w}. \end{aligned}$$

In order to show that $(\mathbf{\Lambda}_0, \mathbf{\Lambda}_1) \in \Omega_c^\varkappa$ implies $\left(\tilde{\mathbf{\Lambda}}_0, \tilde{\mathbf{\Lambda}}_1\right) \in \Omega_c^\varkappa$, we note the following. Assume that $(\mathbf{\Lambda}_0, \mathbf{\Lambda}_1) \in \Omega_c^\varkappa$. Then,

$$\Gamma^*(\mathbf{\Lambda}_0)(\vartheta - \phi) + \Gamma^*(\mathbf{\Lambda}_1)(\varphi - \phi) \leq c(\vartheta - \phi, \varphi - \phi) = c(\vartheta, \varphi) \tag{34}$$

for all $\phi, \vartheta, \varphi \in \mathbb{T}$, where the last equality follows from the assumption that c is shift-invariant. Define

$$\breve{\Phi}_w(\phi) = \Phi_w(-\phi),$$

which satisfies $\int \breve{\Phi}_w(\phi) d\phi \leq 2\pi$ and $\breve{\Phi}_w(\phi) \geq 0$, $\forall \phi \in \mathbb{T}$. Multiplying both sides of (34) with $\breve{\Phi}_w(\phi)$ and integrating with respect to ϕ then yields

$$\frac{1}{2\pi}\left(\Gamma^*(\mathbf{\Lambda}_0) * \breve{\Phi}_w\right)(\vartheta) + \frac{1}{2\pi}\left(\Gamma^*(\mathbf{\Lambda}_1) * \breve{\Phi}_w\right)(\varphi) \leq c(\vartheta, \varphi)$$

for all $\vartheta, \varphi \in \mathbb{T}$, where both sides have been divided by 2π. Similarly,

$$\frac{1}{2\pi}\left(\Gamma^*(\mathbf{\Lambda}_0) * \breve{\Phi}_w\right)(\vartheta) \leq \varkappa, \ \frac{1}{2\pi}\left(\Gamma^*(\mathbf{\Lambda}_1) * \breve{\Phi}_w\right)(\varphi) \leq \varkappa \tag{35}$$

for all $\vartheta, \varphi \in \mathbb{T}$. Further, for any $\vartheta \in \mathbb{T}$, we have

$$\begin{aligned} \frac{1}{2\pi}\left(\Gamma^*(\mathbf{\Lambda}_0) * \breve{\Phi}_w\right)(\vartheta) &= \frac{1}{4\pi^2}\int_{\mathbb{T}} \mathbf{a}(\varphi)^H \mathbf{\Lambda}_0 \mathbf{a}(\varphi) \breve{\Phi}_w(\vartheta - \varphi) d\varphi \\ &= \frac{1}{4\pi^2}\mathrm{tr}\left(\mathbf{\Lambda}_0 \int_{\mathbb{T}} \mathbf{a}(\varphi)\mathbf{a}(\varphi)^H \breve{\Phi}_w(\vartheta - \varphi) d\varphi\right) \\ &= \frac{1}{4\pi^2}\langle \mathbf{\Lambda}_0, \int_{\mathbb{T}} \mathbf{a}(\varphi)\mathbf{a}(\varphi)^H \Phi_w(\varphi - \vartheta) d\varphi \rangle \end{aligned}$$

$$
\begin{aligned}
&= \frac{1}{4\pi^2}\langle \mathbf{\Lambda}_0, \int_{\mathbb{T}} \mathbf{a}(\varphi)\mathbf{a}(\varphi)^H \left(\Phi_w * \delta_\vartheta\right)(\varphi) d\varphi \rangle \\
&= \frac{1}{2\pi}\langle \mathbf{\Lambda}_0, \Gamma\left(\Phi_w * \delta_\vartheta\right) \rangle,
\end{aligned}
$$

where the last equality uses the definition of the operator $\Gamma(\cdot)$. As $\delta_\vartheta(\varphi)$ is a spectrum consistent with the rank-one covariance matrix $\mathbf{R}_\vartheta = \frac{1}{2\pi}\mathbf{a}(\vartheta)\mathbf{a}(\vartheta)^H$, we have, by the properties of the Fourier transform, that $\left(\Phi_w * \delta_\vartheta\right)(\varphi)$ is a spectrum consistent with the covariance matrix $2\pi\mathbf{R}_w \odot \mathbf{R}_\vartheta$, i.e.,

$$
\frac{1}{2\pi}\langle \mathbf{\Lambda}_0, \Gamma\left(\Phi_w * \delta_\vartheta\right) \rangle = \langle \mathbf{\Lambda}_0, \mathbf{R}_w \odot \mathbf{R}_\vartheta \rangle,
$$

and therefore

$$
\begin{aligned}
\frac{1}{2\pi}\left(\Gamma^*(\mathbf{\Lambda}_0) * \breve{\Phi}_w\right)(\vartheta) &= \langle \mathbf{\Lambda}_0, \mathbf{R}_w \odot \mathbf{R}_\vartheta \rangle \\
&= \langle \mathbf{\Lambda}_0 \odot \overline{\mathbf{R}_w}, \mathbf{R}_\vartheta \rangle \\
&= \langle \tilde{\mathbf{\Lambda}}_0, \mathbf{R}_\vartheta \rangle \\
&= \langle \tilde{\mathbf{\Lambda}}_0, \frac{1}{2\pi}\mathbf{a}(\vartheta)\mathbf{a}(\vartheta)^H \rangle \\
&= \frac{1}{2\pi}\mathbf{a}(\vartheta)^H \tilde{\mathbf{\Lambda}}_0 \mathbf{a}(\vartheta) \\
&= \Gamma^*(\tilde{\mathbf{\Lambda}}_0)(\vartheta),
\end{aligned}
$$

where the third equality follows from the definition of $\tilde{\mathbf{\Lambda}}_0$ and the last from the definition of $\Gamma^*(\cdot)$. Using the same reasoning for $\mathbf{\Lambda}_1$, we thus have that $(\mathbf{\Lambda}_0, \mathbf{\Lambda}_1) \in \Omega_c^\varkappa$ implies $\left(\tilde{\mathbf{\Lambda}}_0, \tilde{\mathbf{\Lambda}}_1\right) \in \Omega_c^\varkappa$. Combined, this yields

$$
T_\varkappa(\mathbf{R}_0', \mathbf{R}_1') \leq \max_{(\tilde{\mathbf{\Lambda}}_0, \tilde{\mathbf{\Lambda}}_1) \in \Omega_c^\varkappa} \langle \tilde{\mathbf{\Lambda}}_0, \mathbf{R}_0 \rangle + \langle \tilde{\mathbf{\Lambda}}_1, \mathbf{R}_1 \rangle \; = T_\varkappa(\mathbf{R}_0, \mathbf{R}_1).
$$

□

References

[1] J. P. Burg, "Maximum Entropy Spectral Analysis," in *Proceedings of the 37th Annual International SEG Meeting*, Oklahoma City, OK., 1967.

[2] P. Stoica and R. Moses, *Spectral Analysis of Signals*, Prentice Hall, Upper Saddle River, N.J., 2005.

[3] A. P. Dempster, "Covariance Selection," *Biometrics*, vol. 28, no. 1, pp. 157–175, 1972.

[4] V. Chandrasekaran, P. A. Parrilo, and A. S. Willsky, "Latent variable graphical model selection via convex optimization," *The Annals of Statistics*, vol. 40, no. 4, pp. 1935–1967, 2012.

[5] A. K. Kshirsagar, *Multivariate Analysis*, Dekker, 1972.

[6] A. M. Zoubir, V. Koivunen, Y. Chakhchoukh, and M. Muma, "Robust Estimation in Signal Processing: A Tutorial-Style Treatment of Fundamental Concepts," *IEEE Signal Process. Mag.*, vol. 29, no. 4, pp. 61–80, July 2012.

[7] E. Ollila, D. E. Tyler, V. Koivunen, and H. V. Poor, "Complex Elliptically Symmetric Distributions: Survey, New Results and Applications," *IEEE Trans. Signal Process.*, vol. 60, no. 11, pp. 5597–5625, Nov 2012.

[8] E. Ollila and D. E. Tyler, "Regularized M-Estimators of Scatter Matrix," *IEEE Trans. Signal Process.*, vol. 62, no. 22, pp. 6059–6070, Nov 2014.

[9] A. Wiesel, "Geodesic Convexity and Covariance Estimation," *IEEE Trans. Signal Process.*, vol. 60, no. 23, pp. 6182–6189, Dec 2012.

[10] Y. Sun, P. Babu, and D. P. Palomar, "Robust Estimation of Structured Covariance Matrix for Heavy-Tailed Elliptical Distributions," *IEEE Trans. Signal Process.*, vol. 64, no. 14, pp. 3576–3590, July 2016.

[11] J. P. Burg, D. G. Luenberger, and D. L. Wenger, "Estimation of Structured Covariance Matrices," *Proc. IEEE*, vol. 70, no. 9, pp. 963–974, 1982.

[12] C. L. Byrnes, T. T. Georgiou, and A. Lindquist, "A new approach to spectral estimation: a tunable high-resolution spectral estimator," *IEEE Trans. Signal Process.*, vol. 48, no. 11, pp. 3189–3205, Nov. 2000.

[13] H. Li, P. Stoica, and J. Li, "Computationally efficient maximum likelihood estimation of structured covariance matrices," *IEEE Trans. Signal*, vol. 47, no. 5, pp. 1314–1323, May 1999.

[14] T. T. Georgiou, "Structured covariances and related approximation questions," in *Directions in Mathematical Systems Theory and Optimization*. 2003, pp. 135–140, Springer.

[15] J. Karlsson and P. Enqvist, "Input-to-state covariances for spectral analysis: The biased estimate," in *Proc. of Int. Symp. Mathematical Theory of Networks and Systems*, 2012.

[16] E. Avventi, A. G. Lindquist, and B. Wahlberg, "ARMA Identification of Graphical Models," *IEEE Trans. Autom. Control*, vol. 58, no. 5, pp. 1167–1178, May 2013.

[17] M. Zorzi and R. Sepulchre, "AR identification of latent-variable graphical models," *IEEE Trans. Autom. Control*, vol. 61, no. 9, pp. 2327–2340, Oct 2016.

[18] K. Yamamoto, Y. Chen, L. Ning, T. T. Georgiou, and A. Tannenbaum, "Regularization and Interpolation of Positive Matrices," *arXiv preprint arXiv:1611.07945*, 2016.

[19] F. Barbaresco, "Information geometry of covariance matrix: Cartan-Siegel homogeneous bounded domains, Mostow/Berger fibration and Frechet media," in *Matrix Information Geometry*. 2013, pp. 199–255, Springer.

[20] M. Knott and C. S. Smith, "On the optimal mapping of distributions," *Journal of Optimization Theory and Applications*, vol. 43, no. 1, pp. 39–49, 1984.

[21] L. Ning, X. Jiang, and T. T. Georgiou, "On the Geometry of Covariance Matrices," *IEEE Signal Process. L*, vol. 20, no. 8, pp. 787–790, Aug. 2013.

[22] C. Villani, *Optimal transport: old and new*, Springer Science & Business Media, 2008.

[23] S. Kolouri, S. R. Park, M. Thorpe, D. Slepcev, and G. K. Rohde, "Optimal Mass Transport: Signal processing and machine-learning applications," *IEEE Signal Process. Mag.*, vol. 34, no. 4, pp. 43–59, July 2017.

[24] S. Haker, L. Zhu, A. Tannenbaum, and S. Angenent, "Optimal Mass Transport for Registration and Warping," *International Journal of Computer Vision*, vol. 60, no. 3, pp. 225–240, 2004.

[25] F. De Goes, D. Cohen-Steiner, P. Alliez, and M. Desbrun, "An Optimal Transport Approach to Robust Reconstruction and Simplification of 2D Shapes," *Computer Graphics Forum*, vol. 30, no. 5, pp. 1593–1602, 2011.

[26] B. Engquist and B. D. Froese, "Application of the Wasserstein metric to seismic signals," *Communications in Mathematical Sciences*, vol. 12, no. 5, 2014.

[27] T. T. Georgiou, J. Karlsson, and M. S. Takyar, "Metrics for power spectra: an axiomatic approach," *IEEE Trans. Signal Process.*, vol. 57, no. 3, pp. 859–867, Mar. 2009.

[28] X. Jiang, Z. Q. Luo, and T. T. Georgiou, "Geometric Methods for Spectral Analysis," *IEEE Trans. Signal Process.*, vol. 60, no. 3, pp. 1064–1074, Mar. 2012.

[29] Y. Rubner, C. Tomasi, and L. Guibas, "The Earth Mover's Distance as a Metric for Image Retrieval," *International Journal of Computer Vision*, vol. 40, no. 2, pp. 99–121, 2000.

[30] H. Ling and K. Okada, "An Efficient Earth Mover's Distance Algorithm for Robust Histogram Comparison," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 5, pp. 840–853, 2007.

[31] J. Rabin, G. Peyré, J. Delon, and M. Bernot, "Wasserstein Barycenter and its Application to Texture Mixing," in *International Conference on Scale Space and Variational Methods in Computer Vision*, 2011, pp. 435–446.

[32] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein GAN," *arXiv preprint arXiv:1701.07875*, 2017.

[33] J. Adler, A. Ringh, O. Öktem, and J. Karlsson, "Learning to solve inverse problems using Wasserstein loss," *arXiv preprint arXiv:1710.10898*, 2017.

[34] D. G. Luenberger, *Optimization by Vector Space Methods*, John Wiley and Sons, New York, 1969.

[35] A. Friedman, *Foundations of Modern Analysis*, Courier Corporation, 1970.

[36] D. H. Johnson and D. E. Dudgeon, *Array Signal Processing: Concepts and Techniques*, Prentice Hall, Englewood Cliffs, N.J., 1993.

[37] U. Grenander and G. Szegö, *Toeplitz Forms and Their Applications*, University of California Press, Los Angeles, 1958.

[38] B. N. Bhaskar, G. Tang, and B. Recht, "Atomic Norm Denoising with Applications to Line Spectral Estimation," *IEEE Trans. Signal Process.*, vol. 61, no. 23, pp. 5987 – 5999, July 2013.

[39] F. Elvander, J. Swärd, and A. Jakobsson, "Multi-dimensional grid-less estimation of saturated signals," *Signal Process.*, vol. 145, pp. 37–47, April 2018.

[40] F. Elvander, S. I. Adalbjörnsson, J. Karlsson, and A. Jakobsson, "Using Optimal Transport for Estimating Inharmonic Pitch Signals," in *42nd IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, New Orleans, LA, USA, March 5-9 2017, pp. 331–335.

[41] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004.

[42] D. A. Edwards, "A simple proof in Monge-Kantorovich duality theory," *Studia Mathematica*, vol. 200, no. 1, pp. 67–77, 2010.

[43] S. Jayasumana, R. Hartley, M. Salzmann, H. Li, and M. Harandi, "Kernel Methods on the Riemannian Manifold of Symmetric Positive Definite Matrice," in *2013 IEEE Conference on Computer Vision and Pattern Recognition*, Portland, OR, USA, 23-28 June 2013, pp. 73–80.

[44] J. R. Deller, J. H. L. Hansen, and J. G. Proakis, *Discrete-Time Processing of Speech Signals*, Wiley-IEEE Press, 1999.

[45] B. Dumitrescu, "Trigonometric Polynomials Positive on Frequency Domains and Applications to 2-D FIR Filter Design," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4282–4292, Nov 2006.

[46] B. Dumitrescu, "Optimization of two-dimensional iir filters with nonseparable and separable denominator," *IEEE Trans. Signal Process.*, vol. 53, no. 5, pp. 1768–1777, May 2005.

[47] M. Christensen and A. Jakobsson, *Multi-Pitch Estimation*, Morgan & Claypool, San Rafael, Calif., 2009.

[48] E. Ollila, I. Soloveychik, D. E. Tyler, and A. Wiesel, "Simultaneous penalized M-estimation of covariance matrices using geodesically convex optimization," *arXiv:1608.08126*, 2016.

[49] A. N. Amini and T. T. Georgiou, "Tunable Line Spectral Estimators Based on State-Covariance Subspace Analysis," *IEEE Trans. Signal Process.*, vol. 54, no. 7, pp. 2662–2671, July 2006.

[50] J. Karlsson and T. T. Georgiou, "Uncertainty Bounds for Spectral Estimation," *IEEE Trans. Autom. Control*, vol. 58, no. 7, pp. 1659–1673, July 2013.

B

Paper B

Multi-Marginal Optimal Transport using Partial Information

Filip Elvander[1], Isabel Haasler[2], Andreas Jakobsson[1], and Johan Karlsson[2]

[1] *Centre for Mathematical Sciences, Lund University, Lund, Sweden*
[2] *Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden*

Abstract

During recent decades, there has been a substantial development in optimal mass transport theory and methods. In this work, we consider multi-marginal problems wherein only partial information of each marginal is available, a common setup in many inverse problems in, e.g., remote sensing and imaging. By considering an entropy regularized approximation of the original transport problem, we propose an algorithm corresponding to a block-coordinate ascent of the dual problem, where Newton's algorithm is used to solve the sub-problems. In order to make this computationally tractable for large-scale settings, we utilize the tensor structure that arises in practical problems, allowing for computing projections of the multi-marginal transport plan using only matrix-vector operations of relatively small matrices. As illustrating examples, we apply the resulting method to tracking and barycenter problems in spatial spectral estimation. In particular, we show that the optimal mass transport framework allows for fusing information from different time steps, as well as from different sensor arrays, also when the sensor arrays are not jointly calibrated. Furthermore, we show that by incorporating knowledge of underlying dynamics in tracking scenarios, one may arrive at accurate spectral estimates, as well as faithful reconstructions of spectra corresponding to unobserved time points.

Key words: Optimal mass transport, multi-marginal problems, entropy regularization, spectral estimation, array signal processing, sensor fusion

1 Introduction

The use of mass, or power, distributions for modeling and describing properties of signals and stochastic processes is ubiquitous in the field of signal processing, being fundamental to areas such as spectral estimation, classification, and localization [1]. Spectral representations, detailing the distribution of power over frequency, serve as compact signal descriptions and constitute the foundation for many applications, such as speech enhancement [2], non-invasive estimation [3], spectroscopy [4], as well as radar and sonar [5]. In such applications, one requires a means of comparing and quantifying distances between spectra. Popular choices include the Kullback-Leibler [6] and Itakura-Saito divergences [7], the Rényi entropy [8], and L_p-norms, where the latter is often implicitly utilized through the use of matrix norms applied to the signal covariance matrix [9, 10]. As an alternative to these measures, it has also been proposed to define distances between spectra based on the concept of optimal mass transport [11]. Optimal mass transport (OMT) is concerned with the task of moving one distribution of mass into another, with the description of the movement of mass being referred to as a transport plan. The minimal total cost of rearrangement, i.e., the one associated with the optimal plan, may then be used as a measure of distance between the two distributions [12]. Historically, OMT has been widely used in economics, e.g., for planning and logistics (see [12] for an introduction and an overview of the topic). Recently, there has been a rapid development in the theory and methods for optimal mass transport, and the ideas have attracted considerable attention in various economic and engineering fields, being used, for example, for option pricing [13], color and texture transfer in image processing [14], as well as machine learning [15, 16] and other signal processing tasks [17]. These developments have lead to a mature framework for OMT, with computationally efficient algorithms [18, 19], and constitute a flexible modeling tool that may be used to address many problems in the areas of signal processing and systems theory.

One of the appeals of using OMT for defining distances between distributions lies in its geometric properties. Specifically, as it models transport between two mass distributions, taking place on their domain of definition, the distance is not only related to the point-wise differences between the distributions, but also depends on where the mass is located. As a consequence of this, OMT distances directly reflect proximity of mass on the underlying space [11]. This property becomes very attractive when used in array processing and spatial spectral estimation problems, in which the support of the spectrum directly details the location of sig-

nal emitting sources [1]. For example, OMT has been shown to produce robust and predictable results when applied to direction of arrival (DoA) and localization problems [20, 21]. Also, when computing the distance between two mass distributions, the OMT problem naturally induces a way of interpolating distributions by considering the flow of mass on the underlying space [22, 23]. This procedure has been shown to yield physically meaningful results, e.g., when used for modeling heat diffusion on graphs [24], as well as for interpolating and tracking speech/sound signals [25, 26] and (structured) covariance matrices [27–29].

In this work, we consider multi-marginal OMT problems where the distributions themselves are not directly observed, but manifested through linear measurement equations. We refer to this problem as multi-marginal OMT with partial information. Several problems of interest may be expressed within such a framework, e.g., CT imaging (cf. [30, 31]), ensemble estimation [32], image deblurring [33], spectral analysis [25], and radar imaging problems [21]. As motivating examples, we consider problems from radar imaging, being spatial spectral estimation problems, wherein the mass distributions, i.e., power spectra, are observable from covariance matrices corresponding to a set of sensor arrays. For such problems, typically only the co-located sensors, i.e., sensors in an array, can be processed coherently due to calibration errors between the sensor arrays. Herein, we propose an OMT framework for fusing the information provided by the array covariances in a non-coherent manner. For such scenarios, we demonstrate that multi-marginal OMT constitutes a flexible and robust tool for performing information fusion, illustrating the inherent robustness to spectral perturbations, as caused by, e.g., calibration errors, provided by the use of transport models. In addition to this, we show that the proposed OMT framework may be used to fuse information corresponding to different time points, allowing for modeling the time evolution of spatial spectra. In this case, we demonstrate how the framework allows for incorporating prior knowledge of underlying dynamics and the signal generating mechanism, providing a means of forming accurate reconstructions of target trajectories in spatial estimation problems.

In addition to the modeling framework, one of the main contributions of this work is deriving a computationally efficient method for solving multi-marginal OMT problems with partial information. As the original OMT formulation is posed in an infinite-dimensional function space, we formulate discrete approximations of this problem, allowing for practical implementations. Specifically, one of the main contributions of this work is proposing a generalization of the Sinkhorn

iterations, which may be applied to OMT problems with partial information. In particular, we consider the corresponding dual problem and propose to solve this with a block-coordinate ascent method, wherein each block update is computed using Newton's method. Furthermore, for certain multi-marginal OMT problems, such as the tracking and barycenter problems, we show that the structure of the cost function may be exploited, allowing for solving high-dimensional problems with a large number of marginals. It should be stressed that even though spatial spectral estimation is here used as a motivating and illustrating example, the proposed framework applies to a large class of multi-marginal OMT problems where only partial information of the marginal distributions is available.

This work is based on several of our previous works on tracking and information fusion. In particular, the original tracking formulation is from [32]. This was further developed for radar problems in [20], with efficient computations using pairwise entropy regularization being presented in [21]. In this paper, by extending this work to the multi-marginal setting, we obtain a more flexible framework, with improved computational efficiency and numerical stability.

This paper is structured as follows. In Section 2, we provide background on the problem of optimal mass transport, multi-marginal optimal transport, and partial information. In Section 3, we present the multi-marginal OMT problem with partial information, and describe the use of this formulation in tracking and information fusion. Section 4 presents computational tools for the multi-dimensional OMT problem, which is then extended to problems with partial information. In Section 5, we derive efficient computations for the algorithms in Section 4, by exploiting structures which arise in certain multi-marginal optimal transport problems. In Section 6, we provide numerical examples, illustrating the properties of the proposed framework, whereas Section 7 concludes upon the work.

2 Background: Optimal mass transport

In this section, we provide a brief background on the problem of OMT, together with extensions to settings with underlying dynamics and incomplete information.

2.1 Monge's problem and the Kantorovich relaxation

The problem of OMT, as formulated by Monge, is concerned with finding a mapping between two distributions of mass such that a measure of cost, interpreted as the cost of transportation, is minimized. That is, given two non-negative functions $\Phi_0 \in L_1(\mathcal{X}_0)$ and $\Phi_1 \in L_1(\mathcal{X}_1)$, defined on the spaces $\mathcal{X}_0$ and $\mathcal{X}_1$, respectively, one seeks a function $g : \mathcal{X}_0 \to \mathcal{X}_1$ minimizing the functional [12]

$$\int_{\mathcal{X}_0} c\left(x_0, g(x_0)\right) \Phi_0(x_0) dx_0, \tag{1}$$

subject to the constraint

$$\int_{x_1 \in \mathcal{U}} \Phi_1(x_1) dx_1 = \int_{g(x_0) \in \mathcal{U}} \Phi_0(x_0) dx_0, \tag{2}$$

for all $\mathcal{U} \subset \mathcal{X}_1$. Here, $c : \mathcal{X}_0 \times \mathcal{X}_1 \to \mathbb{R}_+$ denotes the cost of transporting one unit of mass from points in $\mathcal{X}_0$ to points in $\mathcal{X}_1$, yielding that (1) equals the total cost of transporting the distribution Φ_0 to Φ_1 for a given mapping g. The constraint in (2) enforces the condition that g transports Φ_0 to Φ_1, i.e., the mapping g moves the mass in Φ_0 onto Φ_1 without any loss or addition of mass in the transition. This may be expressed using the shorthand notation $g_\# \Phi_0 = \Phi_1$, where $g_\# \Phi_0$ denotes the so-called push-forward measure, i.e.,

$$\int_{x_1 \in \mathcal{U}} \left(g_\# \Phi_0\right)(x_1) dx_1 = \int_{g(x_0) \in \mathcal{U}} \Phi_0(x_0) dx_0, \tag{3}$$

for all $\mathcal{U} \subset \mathcal{X}_1$. The notation thus implies that for any subset $\mathcal{U}$ of $\mathcal{X}_1$, the mass of $g_\# \Phi_0$ on $\mathcal{U}$ is equal to the mass of Φ_0 on the pre-image of $\mathcal{U}$. It directly follows that the total mass of $g_\# \Phi_0$ equals that of Φ_0. Finding such a g that minimizes (1), in general, constitutes a non-convex, non-trivial problem. In contrast, the Kantorovich relaxation of (1) results in a convex problem by replacing the mapping g with an optimal coupling between Φ_0 and Φ_1. This also allows for generalizing the problem to non-negative mass distributions[1], which we denote $\Phi_0 \in \mathcal{M}_+(\mathcal{X}_0)$ and $\Phi_1 \in \mathcal{M}_+(\mathcal{X}_1)$. The set of valid couplings is the set of measures on the product space $\mathcal{X}_0 \times \mathcal{X}_1$ with marginals coinciding with Φ_0 and

[1] Such mass distributions are sometimes referred to as measures and may contain, e.g., Dirac delta functions.

Φ_1, i.e.,

$$\Omega(\Phi_0, \Phi_1) = \Big\{ M \in \mathcal{M}_+(\mathcal{X}_0 \times \mathcal{X}_1) \mid \int_{\mathcal{X}_1} M(\cdot, x_1) dx_1 = \Phi_0, \int_{\mathcal{X}_0} M(x_0, \cdot) dx_0 = \Phi_1 \Big\}. \tag{4}$$

Here, if $M \in \Omega(\Phi_0, \Phi_1)$, then M is referred to as a transport plan, as the function value $M(x_0, x_1)$ may be interpreted as the amount of mass that is transported between x_0 and x_1. Thus, the marginal constraint (4) replaces the condition $g_\# \Phi_0 = \Phi_1$ used in the Monge formulation. It should also be noted that it is assumed that the mass of Φ_0 and Φ_1 are the same, i.e., $\int_{\mathcal{X}_0} \Phi_0(x_0) dx_0 = \int_{\mathcal{X}_1} \Phi_1(x_1) dx_1$, in which case $\Omega(\Phi_0, \Phi_1)$ is always non-empty. The Monge-Kantorovich problem of optimal mass transport may then be stated as

$$\underset{M \in \Omega(\Phi_0, \Phi_1)}{\text{minimize}} \quad \int_{\mathcal{X}_0 \times \mathcal{X}_1} M(x_0, x_1) c(x_0, x_1) dx_0 dx_1, \tag{5}$$

i.e., it seeks the transport plan M minimizing the cost of transportation between Φ_0 and Φ_1, where the point-wise cost of moving a unit mass from $x_0 \in \mathcal{X}_0$ to $x_1 \in \mathcal{X}_1$ is given by $c(x_0, x_1)$. In contrast with (1), the objective function in (5) is linear in M and the optimization problem always has an optimal solution. Further, for $\Phi_0 \in L_1(\mathcal{X}_0)$ and $\Phi_1 \in L_1(\mathcal{X}_1)$, the problems (1) and (5) are equivalent [12].

The resulting minimal objective value of (5) may be used in order to define a notion of distance $S : \mathcal{M}_+(\mathcal{X}_0) \times \mathcal{M}_+(\mathcal{X}_1) \to \mathbb{R}$ according to

$$S(\Phi_0, \Phi_1) = \min_{M \in \Omega(\Phi_0, \Phi_1)} \int_{\mathcal{X}_0 \times \mathcal{X}_1} M(x_0, x_1) c(x_0, x_1) dx_0 dx_1. \tag{6}$$

Recently, the idea of using an OMT-induced concept of proximity has attracted considerable interest in a plethora of modeling applications, such as serving as the learning criterion in machine learning scenarios [16, 34], for performing texture and color transfer [14], dictionary learning [15], as well as being used as an evaluation metric [35]. Also, it has been used for inducing metric structure on the space of power spectra [11], where it was shown that choosing c as a metric on $\mathcal{X} = \mathcal{X}_0 = \mathcal{X}_1$ in general results in S being metric on $\mathcal{M}_+(\mathcal{X})$. Similar efforts have been directed towards defining distances for Toeplitz covariance matrices [28], allowing for formulating a spectral estimation framework with inherent robustness [20, 21]. One of the main motivations for using OMT-based

distances such as (6) is its ability to reflect the geometric properties of the underlying space, $\mathcal{X}$. As a consequence of this, using OMT distances in order to perform interpolation-like tasks, such as, e.g., morphing of images [17], or, recently, for describing the time evolution of a signal on a graph [24], typically yields meaningful and interpretable signal reconstructions. This should be contrasted with the standard Euclidean metric that, when used in the same applications, has been shown to give rise to fade-in fade-out effects, thereby failing to respect underlying assumptions for the signal generating mechanisms, such as, e.g., finite movement speed [11, 24, 28].

2.2 Selection of cost function

It may be noted that the specific choice of cost function c used in the OMT objective in (5) may greatly influence the transport plan M representing the optimal solution. Preferably, c should be chosen as to reflect properties or limitations of the considered application, e.g., in order to ensure expected smooth trajectories of transported mass that may be dictated by physical considerations. In order to allow for incorporating knowledge of the dynamics of the process generating the signal observations, we here utilize the idea to model the underlying dynamics using a linear system which was proposed in [36]. Specifically, let the underlying space be $\mathcal{X} \subset \mathbb{R}^d$, and consider a mass particle with time-varying location, or state, $x(t) \in \mathcal{X}$. Furthermore, assume that the dynamics of the particle may be well modeled by the (continuous time) linear differential equation[2]

$$\dot{x}(t) = Ax(t) + Bu(t), \tag{7}$$

where $x(t) \in \mathcal{X}$, $A \in \mathbb{R}^{d\times d}$, $B \in \mathbb{R}^{d\times f}$, and where $u(t) \in \mathbb{R}^f$ denotes an input signal, steering the evolution of the state $x(t)$. Following the theory from [36], we will herein consider this type of model in order to define a cost function $c : \mathcal{X} \times \mathcal{X} \to \mathbb{R}_+$, describing the cost of transport for a unit mass between two points of $\mathcal{X}$, corresponding to the time instances $t = t_0$ and $t = t_1$. For notational simplicity we here present the case $t_0 = 0$ and $t_1 = 1$. Specifically, we define the cost for transporting mass between $x_0 = x(0)$ to $x_1 = x(1)$ via a

[2] For notational simplicity, we here limit the exposition to the time-invariant case, although this may be generalized to a dynamical system in a straightforward manner.

quadratic optimal control problem, according to

$$\begin{aligned} c(x_0, x_1) = \quad & \min_{u} \quad && \int_0^1 \|u(t)\|_2^2 \, dt \\ & \text{subject to} && \dot{x}(t) = Ax(t) + Bu(t) \\ & && x(0) = x_0, \; x(1) = x_1, \end{aligned} \tag{8}$$

i.e., the cost is equal to the minimal input signal energy required to steer the state from x_0 at $t = 0$ to x_1 at $t = 1$. For this case, the cost may be expressed in closed form as

$$c(x_0, x_1) = (x_1 - \mathcal{A}(1,0)x_0)^T W(1,0)^{-1}(x_1 - \mathcal{A}(1,0)x_0), \tag{9}$$

where $\mathcal{A}(t, \tau) = \text{expm}\,(A(t - \tau))$ is the state transition matrix of (7), with expm$(\cdot)$ denoting the matrix exponential, and where

$$W(t, s) = \int_s^t \mathcal{A}(t, \tau) B B^T \mathcal{A}(t, \tau)^T d\tau$$

is the controllability Gramian. It may be noted that the problem in (8) admits closed form solutions also for the time-varying case (see, e.g., [36]).

Thus, as the cost of transport is dictated by the system dynamics, the optimal transport plan is the coupling $M \in \mathcal{M}_+(\mathcal{X} \times \mathcal{X})$ that provides the mass association most in accordance with the trajectories implied by (7). Furthermore, having access to a transport plan, M, that minimizes the OMT criterion with the dynamic cost in (8) directly allows for computing intermediate mass distributions interpolating the given marginals using the system dynamics in (7); this is termed a displacement interpolation [22]. Specifically, for a particle $x(t)$ such that $x(0) = x_0$ and $x(1) = x_1$, by solving (7) the optimal trajectory is [36, equation (25)]

$$\begin{aligned} \hat{x}(t; x_0, x_1) = \; & \mathcal{A}(t, 1) W(1, t) W(1, 0)^{-1} \mathcal{A}(1, 0) x_0 \\ & + W(t, 0) \mathcal{A}(1, t)^T W(1, 0)^{-1} x_1. \end{aligned}$$

Then, one may construct an interpolating mass distribution, or spectrum, Φ_τ^M, corresponding to time $\tau \in [0, 1]$ and parametrized by the transport plan M, according to

$$\Phi_\tau^M(x) = \int_{\mathcal{X} \times \mathcal{X}} M(x_0, x_1) \delta\left(x - \hat{x}(\tau; x_0, x_1)\right) dx_0 dx_1,$$

where δ denotes the Dirac delta function. That is, $\Phi_\tau^M(x)$ describes the mass located at x at time τ by considering the mass of all pairs (x_0, x_1) such that their trajectories contain the point (x, τ). Also, this framework directly allows for defining extrapolating spectra, i.e., distributions of mass for $\tau > 1$, e.g., by letting the input signal be identically zero for $\tau > 1$, implying that the unforced system dynamics are used to describe the time evolution of the state, and thus the spectrum. It may be noted that the interpolation for Toeplitz covariance matrices presented in [28] may be obtained as a special case of this formulation, using $A = 0$ and $B = I$.

2.3 Partial information

In the OMT formulations in Section 2.1, the input data, i.e., the marginals, used for computing distances and obtaining transport plans, are given. However, in many problems these distributions are only partially known, and are in fact often the quantities of interest to be estimated. In this work, we propose modeling tools that may be applied to several such classes of reconstruction and estimation problems, e.g., inverse problems encountered in imaging, radar, tomography, and signal processing. However, in order to give a concise description, we will herein focus on a few motivating applications in spatial spectral estimation and radar imaging. In such reconstruction problems, one typically seeks to recover a spectrum based on a number of measurements, e.g., estimates of the covariance function of a time series, or estimates of the covariance matrices corresponding to a set of sensor arrays. In this subsection, we describe the basic mathematical model relating the covariance estimates to the power spectra in a localization problem. In the subsequent sections, we will develop theory and formulate OMT problems which may be used to fuse such measurements obtained from different sensors or from different time points in an informed and dependable manner.

2.3.1 Localization and direction of arrival estimation

Consider a scenario in which point-like sources located in a space $\mathcal{X} \subset \mathbb{R}^d$ emit waves impinging on a set of sensors located at $x_k \in \mathbb{R}^d$ for $k = 1, \ldots, q$, with $d \leq 3$. Letting the corresponding sensor signal vector be denoted $z \in \mathbb{C}^q$, one may consider the sensor array covariance matrix, often called the spatial covariance matrix,

$$R = \mathbb{E}\left(zz^H\right).$$

Then, the covariance R may be related to a positive measure on $\mathcal{X}$, denoted $\Phi \in \mathcal{M}_+(\mathcal{X})$, according to $R = \Gamma(\Phi)$, where the linear operator $\Gamma : \mathcal{M}(\mathcal{X}) \to \mathbb{M}^q$ is given by [37]

$$\Gamma(\Phi) = \int_{\mathcal{X}} a(x)\Phi(x)a(x)^H dx, \tag{10}$$

where dx denotes the Lebesgue measure on $\mathcal{X}$, where $\mathbb{M}^q$ denotes the space of Hermitian matrices of size $q \times q$, and where $\mathcal{M}(\mathcal{X})$ denotes the set of signed measures on $\mathcal{X}$. Here, $a : \mathcal{X} \to \mathbb{C}^q$ denotes the array steering vector, i.e., the manifold vector, whose specific functional form depends on the sensor locations, the wavelength, as well as on the propagation properties of this space. For example, the array manifold vector corresponding to spherical wave fronts in $\mathcal{X} \subset \mathbb{R}^d$ may be defined as [38]

$$a(x) = \left(\|x_k - x\|_2^{-(d-1)/2} \exp\left(-2\pi\mathrm{i}\frac{\|x_k - x\|_2}{\xi} \right) \right)_{k=1}^{q}, \tag{11}$$

where $\|\cdot\|_2$ denotes the Euclidean norm, i is the imaginary unit, and ξ is the source signal wavelength. Here, Φ is referred to as a spatial spectrum, as it describes the distribution of mass, or power, on $\mathcal{X}$, with the locations of the signal sources on $\mathcal{X}$ corresponding to the support of Φ.

In this work, we will use the problem of localization, seen as a spectral estimation problem, to illustrate how to utilize the proposed multi-marginal OMT framework for fusing information. Specifically, we will consider the problems of sensor fusion and spectral tracking. Sensor fusion corresponds to combining measurements obtained from separate sensor arrays in forming a single estimate of the spatial spectrum, whereas tracking refers to combining measurements collected at consecutive time points in order to reconstruct the time evolution of a spatial spectrum.

2.3.2 Dynamical models with partially observed states

In the context of using dynamical models such as (7), the partial information contained in the measurements may be related to an output space of lower dimension than the state space $\mathcal{X}$. Specifically, the system description in (7) may be extended with an observation equation according to

$$\begin{aligned} \dot{x}(t) &= Ax(t) + Bu(t) \\ y(t) &= Fx(t), \end{aligned} \tag{12}$$

where $F \in \mathbb{R}^{g \times d}$ is a matrix describing the output of the system, implying that covariances are related to spectra detailing mass distributions on a space $\mathcal{Y} \subset \mathbb{R}^g$, potentially different from $\mathcal{X}$, where $y(t) \in \mathcal{Y}$ is the observed state. Using the system in (12), one may thus incorporate prior knowledge of underlying dynamics, expressed through the state equation, even though one only has access to partial measurements of an observed state. For the corresponding spectra, it holds that if $\Phi \in \mathcal{M}_+(\mathcal{X})$, then $F_\# \Phi \in \mathcal{M}_+(\mathcal{Y})$, where $F_\#$ denotes the push-forward operation, as defined in (3). Note that, using the system in (12), covariance measurements corresponding to observable quantities are related to spectra on the space $\mathcal{Y}$, i.e., to elements of $\mathcal{M}_+(\mathcal{Y})$, and only have an indirect link to spectra on the space $\mathcal{X}$. However, the push-forward operation $F_\#$ allows for expressing the connection between covariance matrices and elements of $\mathcal{M}_+(\mathcal{X})$ by modifying the operators Γ to incorporate $F_\#$ according to

$$\Gamma(\Phi) = \int_{\mathcal{Y}} a(y)\,(F_\# \Phi)\,(y)\, a(y)^H dy. \tag{13}$$

It is worth noting that the domain of definition of the linear operator Γ is still $\mathcal{M}(\mathcal{X})$, whereas the domain of definition for the array steering vectors is $\mathcal{Y}$, reflecting that the covariance measurements are related to elements of $\mathcal{M}_+(\mathcal{Y})$. Also, the transport plan M is still an element of $\mathcal{M}_+(\mathcal{X} \times \mathcal{X})$, allowing for modeling dynamical structures on the space generating data, i.e., $\mathcal{X}$, that are not present in the measurement space $\mathcal{Y}$. Thus, in this context, the notion of partial information is twofold. Firstly, mass distributions are not observed directly, but only through linear measurements on a considerably smaller dimension, i.e., covariances in a finite-dimensional space as opposed to the infinite-dimensional space of measures for spectra. Secondly, these spectra are mappings of spectra that are elements of an even larger space. We will in the next section examine how these ideas, together with the concept of multi-marginal transport plans, can be used in order to arrive at compact formulations for spectral estimation.

3 Multi-marginal optimal mass transport

It is worth noting that the formulation in (5) seeks a transport plan connecting two mass distributions, or marginals, Φ_0 and Φ_1, which may be interpreted as a description of how to morph Φ_0 into Φ_1. A natural extension of this is to consider transport between a larger set of marginals, i.e., a set $\{\Phi_\mathcal{T}\}_{t=0}^{\mathcal{T}}$ for $\mathcal{T} \geq 1$ [39].

To this end, let $\mathbf{M} \in \mathcal{M}_+(\boldsymbol{\mathcal{X}})$, where $\boldsymbol{\mathcal{X}} = \mathcal{X}_0 \times \mathcal{X}_1 \times \cdots \times \mathcal{X}_\mathcal{T}$, and define the set of projection operators $\mathcal{P}_t : \mathcal{M}(\boldsymbol{\mathcal{X}}) \to \mathcal{M}(\mathcal{X}_t)$, for $t = 0, \ldots, \mathcal{T}$, as

$$\mathcal{P}_t(\mathbf{M}) = \int_{\boldsymbol{\mathcal{X}}_{-t}} \mathbf{M}(x_0, \ldots, x_t, \ldots, x_\mathcal{T}) d\mathbf{x}_{-t}, \tag{14}$$

where $d\mathbf{x}_{-t} = dx_0 \ldots dx_{t-1} dx_{t+1} \ldots dx_\mathcal{T}$ and

$$\boldsymbol{\mathcal{X}}_{-t} = \mathcal{X}_0 \times \cdots \times \mathcal{X}_{t-1} \times \mathcal{X}_{t+1} \times \cdots \times \mathcal{X}_\mathcal{T}.$$

It may be noted that one may express marginalization also for the standard two-marginal case using the operator $\mathcal{P}_t$. Then, $\mathbf{M} \in \mathcal{M}_+(\boldsymbol{\mathcal{X}})$ is referred to as a multi-marginal transport plan for the set $\{\Phi_t\}_{t=0}^{\mathcal{T}}$ if $\mathcal{P}_t(\mathbf{M}) = \Phi_t$, for $t = 0, \ldots, \mathcal{T}$. Note here that $\mathbf{M}$ provides a complete description of the association of mass for the set $\{\Phi_t\}_{t=0}^{\mathcal{T}}$, i.e., $\mathbf{M}(x_0, x_1, \ldots, x_\mathcal{T})$ denotes the amount of mass at x_0, corresponding to the marginal Φ_0, that is transported to location x_t, corresponding to marginal Φ_t, for $t = 1, 2, \ldots, \mathcal{T}$. Correspondingly, one may define a cost function $\boldsymbol{\mathcal{C}} : \boldsymbol{\mathcal{X}} \to \mathbb{R}_+$, yielding a generalization of the OMT problem as (see also [39])

$$\begin{aligned} &\underset{\mathbf{M} \in \mathcal{M}_+(\boldsymbol{\mathcal{X}})}{\text{minimize}} \quad \int_{\boldsymbol{\mathcal{X}}} \mathbf{M}(\mathbf{x}) \boldsymbol{\mathcal{C}}(\mathbf{x}) d\mathbf{x} \\ &\text{subject to} \quad \mathcal{P}_t(\mathbf{M}) = \Phi_t, \qquad \text{for } t = 0, \ldots, \mathcal{T}, \end{aligned} \tag{15}$$

where $\mathbf{x} = (x_0, \ldots, x_\mathcal{T})$ and $d\mathbf{x} = dx_0 dx_1 \ldots dx_\mathcal{T}$. Thus, the minimizer of (15) is the transport plan associated with the minimal cost of transportation for the set $\{\Phi_t\}_{t=0}^{\mathcal{T}}$. This type of formulation has previously found application in, e.g., tomography [40] and fluid dynamics [41], with actual numerical approximations and implementations being considered in [30]. In this work, we consider extending this problem to scenarios in which the marginals $\{\Phi_t\}_{t=0}^{\mathcal{T}}$ are not directly observable. Instead, we assume that one has access to partial information in the form of linear mappings of the marginals. In the context of spatial spectral estimation, this corresponds to the spatial covariance matrix, or estimates thereof, i.e., a sequence $\{R_t\}_{t=0}^{\mathcal{T}}$, where $R_t \in \mathbb{M}^{q_t}$ is related to a marginal $\Phi_t \in \mathcal{M}_+(\mathcal{X}_t)$ through the measurement equation $R_t = \Gamma_t(\Phi_t)$. It should be noted that the operators $\Gamma_t : \mathcal{M}(\mathcal{X}_t) \to \mathbb{M}^{q_t}$ may, in general, be different for different indices t.

With this, the problem considered herein may be stated as

$$\begin{aligned} &\underset{\substack{\mathbf{M}\in\mathcal{M}_+(\boldsymbol{\mathcal{X}})\\ \Delta_t\in\mathbb{C}^{q_t\times q_t}}}{\text{minimize}} \quad \int_{\boldsymbol{\mathcal{X}}} \mathbf{M}(\mathbf{x})\boldsymbol{\mathcal{C}}(\mathbf{x})d\mathbf{x} + \sum_{t=0}^{\mathcal{T}} \gamma_t \|\Delta_t\|_{\mathrm{F}}^2 \\ &\text{subject to} \quad \Gamma_t(\mathcal{P}_t(\mathbf{M})) = R_t + \Delta_t \qquad \text{for } t = 0,\dots,\mathcal{T}, \end{aligned} \tag{16}$$

where $\|\cdot\|_{\mathrm{F}}$ denotes the Frobenius norm and where $\gamma_t > 0$, for $t = 0, 1, \dots, \mathcal{T}$, are user-specified weights. Note that the measurement equations have been augmented by error terms $\Delta_t \in \mathbb{C}^{q_t\times q_t}$ in order to allow for noisy covariance matrix estimates R_t. The parameters γ_t then allow for determining a trade-off between how much consideration should be taken to the transport cost and the measurement error, respectively. For simplicity, we here penalize Δ_t by the squared Frobenius norm. In order to compute couplings between marginals corresponding to index pairs (t, s), one requires the corresponding bi-marginal transport plan. Due to the structure of the multi-marginal transport plan $\mathbf{M}$, this is directly computable by means of the bi-marginal projection operator $\mathcal{P}_{t,s} : \mathcal{M}(\boldsymbol{\mathcal{X}}) \to \mathcal{M}(\mathcal{X}_t \times \mathcal{X}_s)$ defined by

$$\mathcal{P}_{t,s}(\mathbf{M}) = \int_{\boldsymbol{\mathcal{X}}_{-t,-s}} \mathbf{M}(x_0,\dots,x_{\mathcal{T}})d\mathbf{x}_{-t,-s}, \tag{17}$$

where $d\mathbf{x}_{-t,-s} = dx_0 \dots dx_{t-1}dx_{t+1} \dots dx_{s-1}dx_{s+1} \dots dx_{\mathcal{T}}$ and

$$\begin{aligned} \boldsymbol{\mathcal{X}}_{-t,-s} = \mathcal{X}_0 \times \cdots \times \mathcal{X}_{s-1} \times \mathcal{X}_{s+1} \times \cdots \times \\ \mathcal{X}_{t-1} \times \mathcal{X}_{t+1} \times \cdots \times \mathcal{X}_{\mathcal{T}}. \end{aligned}$$

This may then be used, e.g., for interpolation between adjacent marginals for which the cost describes a dynamic model as in Section 2.2. In Section 5, we present computationally efficient methods for computing discrete approximations of the projections in (14) and (17). In fact, computing the marginal projections in (14) will constitute an integral component in the algorithm for solving (16), which will be presented in Section 4.

Next, we will see that a number of spectral estimation problems may be cast in the form (16). As will be clear in the following, the only thing distinguishing these spectral estimation problems is the choice of cost function $\boldsymbol{\mathcal{C}}$. For ease of notation, we will in this description often assume that the domains of the marginals are the same, however, this generalizes straightforwardly to general domains.

3.1 Sensor fusion

Consider a scenario in which waves impinge on a set of $J \in \mathbb{N}$ sensor arrays, with corresponding array manifold vectors $a_j : \mathcal{X} \to \mathbb{C}^{q_j}$ and operators $\Gamma_j : \mathcal{M}(\mathcal{X}) \to \mathbb{M}^{q_j}$, for $j = 1, \dots, J$, where it may be noted that the number of sensors, q_j, may differ among the different arrays. Then, a spatial spectrum $\Phi \in \mathcal{M}_+(\mathcal{X})$ gives rise to a set of covariance matrices R_j, for $j = 1, \dots, J$, where $R_j = \Gamma_j(\Phi)$. Thus, as all sensor array measurements have been generated by the same signal, the problem of estimating the spatial spectrum Φ, and thereby the location of the signal sources, corresponds to the inverse problem of finding Φ such that $R_j = \Gamma_j(\Phi)$, for $j = 1, \dots, J$. However, in practice, calibration errors in the sensor arrays, i.e., errors in the operators Γ_j, may result in there existing no single $\Phi \in \mathcal{M}_+(\mathcal{X})$ that is consistent with all array measurements. For scenarios in which only the covariance matrices for the separate sensor arrays are available, for example if these are estimated locally [42, 43], we propose to address this issue using a spectral barycenter formulation. Specifically, we propose to solve

$$\begin{aligned} \underset{\substack{\Phi_j \in \mathcal{M}_+(\mathcal{X}) \\ \Delta_j \in \mathbb{C}^{q_j \times q_j}}}{\text{minimize}} \quad & \sum_{j=1}^{J} S(\Phi_0, \Phi_j) + \gamma_j \left\| \Delta_j \right\|_F^2 \\ \text{subject to} \quad & \Gamma_j(\Phi_j) = R_j + \Delta_j, \; j = 1, \dots, J, \end{aligned} \tag{18}$$

and estimate the spatial spectrum as the minimizer Φ_0. Here, S is the OMT distance defined in (6). Note here that Φ_j, for $j = 1, \dots, J$, are spectra consistent with the set of observations R_j, whereas Φ_0 is the spectrum closest, in the OMT sense, to this set of spectra, i.e., Φ_0 is the spectral barycenter. It may be noted that the formulation in (18) induces robustness to the estimate Φ_0 as small spatial perturbations caused by, e.g., calibration errors, do not result in large values of the distance measure S. One may formulate the problem in (18) as a multi-marginal problem on the form (16) by specifying the multi-marginal cost function $\boldsymbol{\mathcal{C}}$ as

$$\boldsymbol{\mathcal{C}}(\mathbf{x}) = \sum_{j=1}^{J} c(x_0, x_j), \tag{19}$$

where $c : \mathcal{X} \times \mathcal{X} \to \mathbb{R}_+$ is the pair-wise cost function in the definition of S. With this, the barycenter problem may be stated as

$$\begin{aligned} \underset{\substack{\mathbf{M}\in\mathcal{M}_+(\boldsymbol{\mathcal{X}}) \\ \Delta_j\in\mathbb{C}^{q_j\times q_j}}}{\text{minimize}} \quad & \int_{\boldsymbol{\mathcal{X}}} \mathbf{M}(\mathbf{x}) \left(\sum_{j=1}^{J} c(x_0, x_j) \right) d\mathbf{x} + \sum_{j=1}^{J} \gamma_j \|\Delta_j\|_{\mathrm{F}}^2 \\ \text{subject to} \quad & \Gamma_j(\mathcal{P}_j(\mathbf{M})) = R_j + \Delta_j \qquad \text{for } j = 1, \ldots, J. \end{aligned} \tag{20}$$

To see that (19) indeed induces a barycenter solution when used in (20), note that the cost function $\boldsymbol{\mathcal{C}}$ is structured as to penalize transport between a central marginal, corresponding to the index $j = 0$, and all other marginals, corresponding to indices $j = 1, \ldots, J$. As only marginals with indices $j = 1, \ldots, J$ correspond to actual measurements, this implies that the marginal $t = 0$, which is not constrained to be related to any measurements, should be close, in the OMT sense, to the entire ensemble of marginals. Note here that the barycenter does not appear as an explicit variable, but is instead given by the projection of $\mathbf{M}$ on the zeroth marginal, i.e., $\Phi_0 = \mathcal{P}_0(\mathbf{M})$. It may be noted that, in general, one may utilize different OMT distances S_j for different arrays, e.g., one may replace the sums in (18) and (19) with a weighted sum to get a weighted barycenter.

3.2 Tracking

One may also consider a scenario in which a single sensor array[3] is used to observe the evolution of a signal over time. Specifically, assume that estimates of the array covariance matrix are obtained at $\mathcal{T}+1$ time instances corresponding to the nominal times $t = 0, 1, \ldots, \mathcal{T}$, giving rise to the sequence $\{R_t\}_{t=0}^{\mathcal{T}}$. Under the assumption that the corresponding spatial spectrum should evolve smoothly, implying signal sources moving at finite speed, one may estimate the corresponding sequence of spectra, $\{\Phi_t\}_{t=0}^{\mathcal{T}}$, using the tracking formulation

$$\begin{aligned} \underset{\substack{\Phi_t\in\mathcal{M}_+(\mathcal{X}) \\ \Delta_t\in\mathbb{C}^{q\times q}}}{\text{minimize}} \quad & \sum_{t=1}^{T} S_t(\Phi_{t-1}, \Phi_t) + \gamma_t \|\Delta_t\|_F^2 \\ \text{subject to} \quad & \Gamma(\Phi_t) = R_t + \Delta_t, \ t = 0, \ldots, \mathcal{T}, \end{aligned} \tag{21}$$

i.e., the sequence $\{\Phi_t\}_{t=0}^{\mathcal{T}}$ minimizes a sequential OMT distance, while interpolating the (de-noised) measurement sequence $\{R_t\}_{t=0}^{\mathcal{T}}$. It may here be noted that

[3] We present a formulation for handling multi-array cases in the next section.

the OMT distance, S_t, is indexed by t, as to express that the underlying cost of transport may change depending on the time. Correspondingly, a multi-marginal formulation may be obtained by considering a cost function $\boldsymbol{\mathcal{C}}$ such that

$$\boldsymbol{\mathcal{C}}(\mathbf{x}) = \sum_{t=1}^{\mathcal{T}} c_t(x_{t-1}, x_t), \tag{22}$$

with the index t expressing time-dependence, yielding

$$\begin{aligned} \underset{\substack{\mathbf{M}\in\mathcal{M}_+(\boldsymbol{\mathcal{X}}) \\ \Delta_t\in\mathbb{C}^{q\times q}}}{\text{minimize}} \quad & \int_{\boldsymbol{\mathcal{X}}} \mathbf{M}(\mathbf{x}) \left(\sum_{t=1}^{\mathcal{T}} c_t(x_{t-1}, x_t) \right) d\mathbf{x} + \sum_{t=0}^{\mathcal{T}} \gamma_t \|\Delta_t\|_{\mathrm{F}}^2 \\ \text{subject to} \quad & \Gamma(\mathcal{P}_t(\mathbf{M})) = R_t + \Delta_t \qquad \text{for } t = 0, \ldots, \mathcal{T}. \end{aligned} \tag{23}$$

Here, the estimated spectral sequence is given by the marginals of $\mathbf{M}$, i.e., $\Phi_t = \mathcal{P}_t(\mathbf{M})$, for $t = 0, 1, \ldots, \mathcal{T}$. It is worth noting that the cost function $\boldsymbol{\mathcal{C}}$ penalizes transport between consecutive marginals, i.e., between pairs corresponding to indices $(t-1, t)$, for $t = 1, \ldots, \mathcal{T}$. This allows for exploiting assumptions of an underlying dynamic model, expressed in c_t in (22), as described in Section 2.2. Specifically, let the space on which the spatial spectra are defined be denoted $\mathcal{Y}$, and let $\mathcal{X}$ correspond to a larger space, allowing for expressing the system's dynamics. Using the system description in (12), the cost function c_t may be defined according to (8), and the operators Γ may be extended to include the push-forward operation, as detailed in (13). The sequence of spatial spectra may then be obtained as $\Phi_t = F_{\#}\mathcal{P}_t(\mathbf{M})$. We will in Section 6 examine the implications of using a dynamical model as opposed to a static one when performing OMT-based spectral tracking.

Example 1. Consider a DoA estimation scenario in $\mathbb{R}^2$, in which a sequence of array covariance matrices is available, where one in addition to the direction of arrival, $\vartheta \in [-\pi, \pi)$ aims to model also the velocity of the targets, in order to enforce smooth trajectories. Letting $v(t) \in \mathbb{R}$ denote the angular speed, i.e., $\dot{\vartheta}(t) = v(t)$, one may thus select $\mathcal{Y} = [-\pi, \pi)$ and $\mathcal{X} = [-\pi, \pi) \times \mathbb{R}$. In order to impose the assumption of smooth trajectories, one may then enforce finite acceleration by, in the state space representation in (12), using the matrices

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \; B = \begin{bmatrix} 0 & 1 \end{bmatrix}^T, \; F = \begin{bmatrix} 1 & 0 \end{bmatrix}.$$

It may be noted that the pair (A, B) enforces that the angular position can only be influenced through its angular velocity, i.e., through acceleration, whereas the matrix F expresses the fact that only the angular position is manifested in the array covariance matrix. The corresponding relation between spectra on $\mathcal{X}$ and $\mathcal{Y}$, denoted $\Phi_t^{\mathcal{X}}$ and $\Phi_t^{\mathcal{Y}}$, respectively, may then be expressed as

$$\Phi_t^{\mathcal{Y}}(\vartheta) = \left(F_{\#}\Phi_t^{\mathcal{X}}\right)(\vartheta) = \int_{\mathbb{R}} \Phi_t^{\mathcal{X}}(\vartheta, v)dv,$$

for $\vartheta \in [-\pi, \pi)$. The cost function c implied by the LQ control signal u is then given by (see, e.g., [32])

$$c(x_0, x_1) = \left(x_1 - \operatorname{expm}(A)x_0\right)^T W(1,0)^{-1}\left(x_1 - \operatorname{expm}(A)x_0\right)$$

with

$$\operatorname{expm}(A(t-\tau)) = \begin{bmatrix} 1 & t-\tau \\ 0 & 1 \end{bmatrix}, \quad W(1,0) = \begin{bmatrix} 1/3 & 1/2 \\ 1/2 & 1 \end{bmatrix}.$$

□

3.3 Combining tracking and barycenters

In addition to the problems described in Sections 3.1 and 3.2, one may consider a combination of the two, i.e., the tracking of a spectral barycenter over time. This then allows for addressing issues of resolving spatial ambiguities, as well as promoting smooth spectral trajectories over time as to, e.g., induce robustness to noisy measurements. Therefore, consider J sensor arrays observing a signal during $\mathcal{T}+1$ time instances, giving rise to a set of covariance matrices $R_t^{(j)}$, for $j = 1, \ldots, J$ and $t = 0, \ldots, \mathcal{T}$, where the subscript indicates time, and the superscript corresponds to the array index. Thus, we seek to estimate a sequence of spectral barycenters, $\{\Phi_t\}_{t=0}^{\mathcal{T}}$, given covariance matrix estimates $R_t^{(j)}$, for $t = 0, 1, \ldots, \mathcal{T}$ and $j = 1, 2, \ldots, J$. In order to formulate the barycenter tracking problem, one may consider two separate types of transport costs; S_t detailing the OMT distance between consecutive barycenters, and $\tilde{S}$ describing the distances between barycenters and spectra consistent with the array measurements. For example, the distance S_t may reflect underlying dynamics whereas $\tilde{S}$ may correspond to a static OMT problem. Correspondingly, one may let the

barycenter spectra be elements of $\mathcal{M}_+(\mathcal{X})$, where the space $\mathcal{X}$ may detail the system dynamics, whereas the consistent spectra are elements of $\mathcal{M}_+(\mathcal{Y})$, where $\mathcal{Y}$ is the measurement space.

With this, the barycenter tracking problem may be formulated as

$$\begin{aligned} &\underset{\substack{\Phi_t \in \mathcal{M}_+(\mathcal{X})\,,\ \Phi_t^{(j)} \in \mathcal{M}_+(\mathcal{Y}) \\ \Delta_t^{(j)} \in \mathbb{C}^{q_j \times q_j}}}{\text{minimize}} && \mathbf{S}(\mathbf{\Phi}) + \alpha \tilde{\mathbf{S}}(\mathbf{\Phi}) \\ &\text{subject to} && I^{(j)}\left(\Phi_t^{(j)}\right) = R_t^{(j)} + \Delta_t^{(j)}, \end{aligned} \tag{24}$$

for $t = 0, \ldots, \mathcal{T}$, and $j = 1, \ldots, J$, where

$$\begin{aligned} \mathbf{S}(\mathbf{\Phi}) &= \sum_{t=1}^{\mathcal{T}} S_t\left(\Phi_{t-1}, \Phi_t\right) \\ \tilde{\mathbf{S}}(\mathbf{\Phi}) &= \sum_{t=0}^{\mathcal{T}} \sum_{j=1}^{J} \left(\tilde{S}\left(\Phi_t^{(j)}, F_{\#}\Phi_t\right) + \gamma_t^{(j)} \left\| \Delta_t^{(j)} \right\|_F^2 \right), \end{aligned}$$

with $\alpha > 0$ being a user-specified parameter allowing for deciding a trade-off between the smoothness of the tracking and the distance between a barycenter Φ_t and its corresponding consistent spectra $\Phi_t^{(j)}$, for $t = 0, 1, \ldots, \mathcal{T}$ and $j = 1, 2, \ldots, J$. It may here be noted the latter distance is stated in terms of the push-forward of the barycenter spectrum, thereby allowing for some states to be measured (e.g., position) whereas others are not (e.g., velocity). In order to arrive at a multi-marginal formulation on the form (16), one may define the cost function $\boldsymbol{\mathcal{C}}$ as

$$\boldsymbol{\mathcal{C}}(\mathbf{x}) = \sum_{t=1}^{\mathcal{T}} c_t(x_{t-1,0}, x_{t,0}) + \alpha \sum_{t=0}^{\mathcal{T}} \sum_{j=1}^{J} \tilde{c}(x_{t,0}, x_{t,j}), \tag{25}$$

where $\mathbf{x} \in \boldsymbol{\mathcal{X}}^J = \boldsymbol{\mathcal{X}} \times \boldsymbol{\mathcal{X}} \times \ldots \times \boldsymbol{\mathcal{X}}$, which is the direct product of $\mathcal{T} + 1$ instances of $\boldsymbol{\mathcal{X}}$, each corresponding to a barycenter problem at one time point. That is, this space is a product of the form $\boldsymbol{\mathcal{X}} = \mathcal{X} \times \mathcal{Y} \times \cdots \times \mathcal{Y}$, with one instance of $\mathcal{X}$ and J instances of $\mathcal{Y}$. Here, $\mathbf{x} = (\mathbf{x}_0, \mathbf{x}_1, \ldots, \mathbf{x}_{\mathcal{T}})$, where

$$\mathbf{x}_t = (x_{t,0}, x_{t,1}, \ldots, x_{t,J}),$$

where the first index indicate the time t, for $t = 0, 1, \ldots, \mathcal{T}$, and the second index indicates the array j, for $j = 0, 1, \ldots, J$, with 0 corresponding to the barycenter. Here, c_t denotes the cost of transport for the tracking, whereas $\tilde{c}$ defines the cost of transportation between the barycenters and the marginal spectra. With this setup, one may formulate the multi-marginal problem as

$$\begin{aligned} \underset{\substack{\mathbf{M}\in\mathcal{M}_+(\boldsymbol{\mathcal{X}}^J)\\ \Delta_t^{(j)}\in\mathbb{C}^{q_j\times q_j}}}{\text{minimize}} \quad & \int_{\boldsymbol{\mathcal{X}}^J} \mathbf{M}(\mathbf{x})\boldsymbol{\mathcal{C}}(\mathbf{x})d\mathbf{x} + \alpha\sum_{t=0}^{\mathcal{T}}\sum_{j=1}^{J}\gamma_t^{(j)}\|\Delta_t^{(j)}\|_{\mathrm{F}}^2 \\ \text{subject to} \quad & I^{(j)}(\mathcal{P}_{(t,j)}(\mathbf{M})) = R_t^{(j)} + \Delta_t^{(j)}, \end{aligned} \tag{26}$$

for $t = 0, \ldots, \mathcal{T}$, and $j = 1, \ldots, J$. Here, the operator $\mathcal{P}_{(t,j)}$ projects onto the marginal (t, j) of $\mathbf{M}$ in direct analogy with (14), which for each time t corresponds to a spectrum consistent with the measurements of array j, for $j = 1, 2, \ldots, J$. The corresponding sequence of barycenters may then be constructed as $\Phi_t = \mathcal{P}_{(t,0)}(\mathbf{M})$, for $t = 0, 1, \ldots, \mathcal{T}$.

Having shown that the multi-marginal OMT problem with partial information in (16) may be used to formulate these spectral estimators, we proceed to present the main results of this work, i.e., how to construct discrete approximations of (16) that allow for computationally efficient solvers. Numerical examples that illustrate the behavior of (16) in spatial spectral estimation scenarios are then provided in Section 6.

4 Solving multi-marginal OMT problems with partial information of marginals

The multi-marginal OMT problem with partial information in (16) is an infinite-dimensional problem, as the multi-marginal transport plan $\mathbf{M}$ is an element of the function space $\mathcal{M}_+(\boldsymbol{\mathcal{X}})$. Fortunately, one may in practical implementations approximate (16) by a discrete counterpart. In this section, we present a computationally efficient solution algorithm for such a problem. Starting by describing a discretization of OMT problems with full information, i.e., problems of the form (15), we derive an algorithm for solving discrete OMT problems with partial information. We show that the critical component of the resulting algorithms is to compute the projections in (14). Section 5 presents several cases for which the projections in (14) can be computed efficiently.

4.1 Entropy regularized OMT and Sinkhorn iterations

In order to discretize the OMT problem in (5), one may construct a grid consisting of n grid points $X = \{\mathrm{x}_1, \mathrm{x}_2, \ldots, \mathrm{x}_n\}$ on the support of the marginals. The discretized marginals are then represented by vectors $\Phi_t \in \mathbb{R}^n$ for $t = 0, \ldots, \mathcal{T}$, with $\mathcal{T} \geq 1$. For ease of notation, we will herein assume that all marginals are discretized using the same number of points. However, this, as well as all presented results, generalizes straightforwardly to the case when the number of discretization points differ among the marginals. In the discretized formulation, the cost function c is represented by a $(\mathcal{T}+1)$-dimensional tensor $\mathbf{C} \in \mathbb{R}^{n^{\mathcal{T}+1}}$, where the element $\mathbf{C}_{i_0,i_1,\ldots,i_\mathcal{T}} = \boldsymbol{\mathcal{C}}(\mathrm{x}_{i_0}, \mathrm{x}_{i_1}, \ldots, \mathrm{x}_{i_\mathcal{T}})$ denotes the transport cost associated with the tuple $(\mathrm{x}_{i_0}, \mathrm{x}_{i_1}, \ldots, \mathrm{x}_{i_\mathcal{T}})$. Correspondingly, the transport plan may be represented by a tensor $\mathbf{M} \in \mathbb{R}_+^{n^{\mathcal{T}+1}}$, whose elements $\mathbf{M}_{i_0,i_1,\ldots,i_\mathcal{T}}$ denote the amount of mass associated with the tuple $(\mathrm{x}_{i_0}, \mathrm{x}_{i_1}, \ldots, \mathrm{x}_{i_\mathcal{T}})$. As a discrete analog to the projection in (14), let $P_t : \mathbb{R}^{n^{\mathcal{T}+1}} \to \mathbb{R}^n$ be the projection on the t-th marginal of a tensor $\mathbf{M}$, which is computed by summing over all modes of $\mathbf{M}$ except t, i.e., the j-th element of the vector $P_t(\mathbf{M}) \in \mathbb{R}^n$ is given by

$$P_t(\mathbf{M})_j = \sum_{\substack{i_0,\ldots,i_{t-1},\\ i_{t+1},\ldots,i_\mathcal{T}}} \mathbf{M}_{i_0,\ldots,i_{t-1},j,i_{t+1},\ldots,i_\mathcal{T}}.$$

Similarly, in analog to the bi-marginal projection in (17), let $P_{t_1,t_2}(\mathbf{M}) \in \mathbb{R}^{n\times n}$ denote the projection of $\mathbf{M}$ on the two joint marginals t_1 and t_2, i.e.,

$$P_{t_1,t_2}(\mathbf{M})_{j,\ell} = \sum_{\substack{i_t:\\ t\in\{0,\ldots,\mathcal{T}\}\setminus\{t_1,t_2\}}} \mathbf{M}_{i_0,\ldots,i_{t_1-1},j,i_{t_1+1},\ldots,i_{t_2-1},\ell,i_{t_2+1},\ldots,i_\mathcal{T}}.$$

With this, the discretized multi-marginal OMT problem is given by

$$\begin{aligned} \underset{\mathbf{M}\in\mathbb{R}^{n^{\mathcal{T}+1}}}{\text{minimize}} \quad & \langle \mathbf{C}, \mathbf{M} \rangle \\ \text{subject to} \quad & P_t(\mathbf{M}) = \Phi_t, \quad t = 0, \ldots, \mathcal{T}, \end{aligned} \tag{27}$$

where $\langle \mathbf{C}, \mathbf{M} \rangle \triangleq \sum_{i_0,\ldots,i_\mathcal{T}} \mathbf{M}_{i_0,\ldots,i_\mathcal{T}} \mathbf{C}_{i_0,\ldots,i_\mathcal{T}}$. Although (27) is a linear program, finding an optimal transport plan can be computationally cumbersome due to the large number of variables. In particular, the classical bi-marginal OMT problem with a grid of size n yields a discrete problem with n^2 variables. It may be noted

that standard linear programming methods require $\mathcal{O}(n^3 \log(n))$ operations to find the optimum for such a problem [44], proving infeasible even for moderate grid sizes.

One way to bypass the expensive computations inherent in finding the exact transport plan is to regularize the OMT problem, yielding an approximation of (27). To this end, it was first suggested in [18] to introduce an entropy term to the objective function. In this setting, the problem in (27) is modified to

$$\begin{aligned} \underset{\mathbf{M}\in\mathbb{R}^{n^{\mathcal{T}+1}}}{\text{minimize}} \quad & \langle \mathbf{C}, \mathbf{M} \rangle + \varepsilon \mathcal{D}(\mathbf{M}) \\ \text{subject to} \quad & P_t(\mathbf{M}) = \Phi_t, \quad t = 0, \dots, \mathcal{T}, \end{aligned} \tag{28}$$

where

$$\mathcal{D}(\mathbf{M}) \triangleq \sum_{i_0,\dots,i_{\mathcal{T}}} \left(\mathbf{M}_{i_0,\dots,i_{\mathcal{T}}} \log \left(\mathbf{M}_{i_0,\dots,i_{\mathcal{T}}} \right) - \mathbf{M}_{i_0,\dots,i_{\mathcal{T}}} + 1 \right),$$

and $\varepsilon > 0$ is a small regularization parameter. The entropy regularized OMT problem is related to the Schrödinger bridge problem, which studies the most likely evolution of a particle cloud observed at two different time instances [45, 46]. Note that due to the entropy term, the optimal transport plan $\mathbf{M}$ corresponding to problem (28) always has full support, whereas the solution to (27) may be sparse. However, in contrast to (27), the regularized problem (28) is strictly convex and thus always guarantees a unique solution. Furthermore, in [18], it was shown for the bi-marginal setting that, as $\varepsilon \to 0$, the solution of (28) converges to the solution of (27) with maximal entropy. More importantly, though, the formulation in (28) allows for deriving an efficient method to find the (approximate) optimal transport plan $\mathbf{M}$, as it implies a low-dimensional representation requiring only $(\mathcal{T}+1)n$ variables instead of the original $n^{\mathcal{T}+1}$ variables.

Specifically, in the original bi-marginal case, where the mass transport plan and cost tensor are given by matrices $M, C \in \mathbb{R}^{n\times n}$, the regularized problem allows for expressing the solution M through diagonal scaling of an $n \times n$ constant matrix. To see this, consider the Lagrange function of (28) with $\mathcal{T} = 1$ and dual variables $\lambda_0, \lambda_1 \in \mathbb{R}^n$, i.e.,

$$\begin{aligned} L(M, \lambda_0, \lambda_1) = \operatorname{tr}\left(C^T M \right) + \varepsilon \mathcal{D}(M) \\ + \lambda_0^T (\Phi_0 - M\mathbf{1}) + \lambda_1^T (\Phi_1 - M^T \mathbf{1}). \end{aligned} \tag{29}$$

For fixed dual variables, the minimum of (29) with respect to M is attained when the gradient with respect to the matrix entries m_{ij} vanishes, yielding a solution of the form

$$M = \operatorname{diag}(u_0) K \operatorname{diag}(u_1), \tag{30}$$

where $K = \exp(-C/\varepsilon)$, $u_0 = \exp(\lambda_0/\varepsilon), u_1 = \exp(\lambda_1/\varepsilon)$, and where $\exp(\cdot)$ denotes element-wise application of the exponential function. Then, as the matrix K has strictly positive elements, it follows from Sinkhorn's theorem [47] that there is a unique matrix M on the form (30) with prescribed strictly positive row and column sums Φ_0 and Φ_1. Moreover, the two positive vectors u_0 and u_1 are unique up to multiplication with a scalar and may be found via Sinkhorn iterations, i.e., by iteratively updating u_0 and u_1 as to satisfy the marginal constraints. Specifically, the Sinkhorn iterations are given by [18]

$$\begin{aligned} u_0 &= \Phi_0 ./ (K u_1) \\ u_1 &= \Phi_1 ./ (K^T u_0), \end{aligned} \tag{31}$$

where $./$ denotes elementwise division. These iterations converge linearly [48], with the computational bottleneck of the scheme being the two matrix-vector multiplications. The Sinkhorn iterations thus provide an efficient technique for finding approximate solutions to OMT problems.

In direct analog to the bi-marginal case, by considering the Lagrangian relaxation of (28), one finds that the transport tensor may be represented by "diagonal" scaling of a constant tensor. Specifically, one may write $\mathbf{M} = \mathbf{K} \odot \mathbf{U}$, where $\odot$ denotes elementwise multiplication, for the two tensors $\mathbf{K}, \mathbf{U} \in \mathbb{R}^{n^{\mathcal{T}+1}}$, given by $\mathbf{K} = \exp(-\mathbf{C}/\varepsilon)$ and

$$\mathbf{U} = (u_0 \otimes u_1 \otimes \cdots \otimes u_{\mathcal{T}}), \quad \iff \mathbf{U}_{i_0,\ldots,i_{\mathcal{T}}} = \prod_{t=0}^{\mathcal{T}} (u_t)_{i_t},$$

for a set of vectors $u_0, \ldots, u_{\mathcal{T}} \in \mathbb{R}^n$. It may be noted that this representation for the mass transfer tensor $\mathbf{M}$ is a direct generalization of the representation in the standard bi-marginal OMT theory, where $\mathbf{M}$, $\mathbf{K}$, and $\mathbf{U} = u_0 \cdot u_1^T$ are matrices, and

$$\mathbf{M} = \mathbf{K} \odot \mathbf{U} = \mathbf{K} \odot (u_0 \cdot u_1^T) = \operatorname{diag}(u_0) \mathbf{K} \operatorname{diag}(u_1),$$

Algorithm 1 Sinkhorn method for the multi-marginal optimal mass transport problem [30].

Given: Initial guess $u_t > 0$, for $t = 0, \dots, \mathcal{T}$; starting point t
while Sinkhorn not converged **do**
  $u_t \leftarrow u_t \odot \Phi_t ./ P_t(\mathbf{K} \odot \mathbf{U})$,
  $t \leftarrow t + 1 \pmod{\mathcal{T} + 1}$
end while
return $\mathbf{M} = \mathbf{K} \odot \mathbf{U}$

as in (30). Thus, the problem in (28) may be reduced to finding $\mathcal{T} + 1$ scaling vectors in $\mathbb{R}^n$ instead of directly optimizing over the complete mass tensor $\mathbf{M} \in \mathbb{R}^{n^{\mathcal{T}+1}}$. Interestingly, the scaling vectors $u_0, \dots, u_{\mathcal{T}}$ are unique up to a scalar for any $\mathcal{T} \geq 1$ [48]. Furthermore, they may be determined via a Sinkhorn iteration scheme similar to (31), as presented in [30] for multi-marginal OMT problems. For completeness, we state the result here: given an initial set of positive vectors $u_t \in \mathbb{R}^n$, for $t = 0, 1, \dots, \mathcal{T} + 1$, the Sinkhorn method (Algorithm 1) is to iteratively update according to

$$u_t \leftarrow u_t \odot \Phi_t ./ P_t(\mathbf{K} \odot \mathbf{U}), \quad t = 0, 1, \dots, \mathcal{T}. \tag{32}$$

It may be noted that also this scheme reduces to the standard Sinkhorn iterations (31) for the case $\mathcal{T} = 1$, i.e., a bi-marginal problem. Parallels between the Sinkhorn iteration scheme and established algorithms for convex optimization problems have been explored in several works. For instance, it has been shown that Sinkhorn iterations correspond to iterative Bregman projections [30]. In the bi-marginal case they have been derived as a block-coordinate ascent in the dual formulation of (28) [31]. In this work, we utilize the latter result in order to develop new Sinkhorn-type schemes for OMT problems with partial information of the marginals, which the problem with full information in (28) is a special case of.

4.2 Method for partial information of marginals

In this section, we derive an efficient algorithm for solving discrete OMT problems where only partial information of the marginals are available, i.e., formulating discrete approximations of the problem in (16). In this setting, one does

not have access to the marginal vectors $\Phi_t \in \mathbb{R}^n$ directly, but instead has measurements of the form $r = G\Phi$, where $r \in \mathbb{R}^m$, with $m < n$, is a vector of available information, and where $G \in \mathbb{R}^{m\times n}$ represents a mapping from the full state information to the partial information. For example, G may represent a discretization of the push-forward operator $F_\#$ in the case of an OMT problem with dynamics detailed by (12), or of the operator Γ, mapping spectra to covariance matrices. Correspondingly, r may be a discretization of the push-forward measure, or the (vectorized) covariance matrix, respectively. Note that all quantities in this section are assumed to be real-valued. As the considered covariance matrices R are, in general, complex, we arrive at the equivalent real-valued problems by constructing the corresponding information vectors r as

$$r = \begin{bmatrix} \mathfrak{Re}\left(\mathrm{vec}(R)\right)^T & \mathfrak{Im}\left(\mathrm{vec}(R)\right)^T \end{bmatrix}^T, \tag{33}$$

where $\mathfrak{Re}(\cdot)$ and $\mathfrak{Im}(\cdot)$ denote the real and imaginary parts, respectively. The discretizations of the operators Γ, i.e., G, are then structured as to be consistent with this construction. It should be noted that the solution algorithm presented in this section coincides with the multi-marginal Sinkhorn iterations in (32) for the special case of full information (corresponding to letting G be the identity operator) and noiseless observations of the distribution marginals ($\gamma_t = \infty$).

Consider the multi-marginal OMT problem in (16), where only noisy and partial measurements of the marginals are available. A direct discretization of this problem, with added entropy regularization, may be expressed as

$$\underset{\mathbf{M},\Delta_t}{\text{minimize}} \quad \langle \mathbf{C}, \mathbf{M}\rangle + \varepsilon\mathcal{D}(\mathbf{M}) + \sum_t \gamma_t \|\Delta_t\|_2^2 \tag{34a}$$

$$\text{subject to} \quad G_t P_t(\mathbf{M}) = r_t + \Delta_t \quad \text{for } t = 0, \ldots, \mathcal{T}, \tag{34b}$$

where $G_t \in \mathbb{R}^{m_t\times n}$ is a mapping from the full state information to the partial information $r_t \in \mathbb{R}^{m_t}$, as described above. Also, to allow for noisy measurements r_t, the problem is augmented by perturbation vectors $\Delta_t \in \mathbb{R}^{m_t}$ by direct analog to the matrix perturbations of (16), which are penalized by the squared ℓ_2-norm scaled by the penalty parameters $\gamma_t > 0$. It may be noted that due to the presence of the entropy term, $\mathcal{D}$, the problem in (34) is strictly convex. Also, for finite γ_t, a solution always exists due to the introduction of the perturbation vectors Δ_t. Altogether, this implies that for any $\varepsilon > 0$ and $\gamma_t < \infty$, the optimization problem (34) has a unique solution. The following proposition characterizes the

solution $\mathbf{M}$ in terms of the variables of the Lagrange dual problem, yielding the dual formulation.

Proposition 1. *The optimal solution to* (34) *may be expressed as* $\mathbf{M} = \mathbf{K} \odot \mathbf{U}$*, for two tensors* $\mathbf{K}, \mathbf{U} \in \mathbb{R}^{n^{\mathcal{T}+1}}$*, given by* $\mathbf{K} = \exp(-\mathbf{C}/\varepsilon)$*, and*

$$\mathbf{U} = u_0 \otimes u_1 \otimes \cdots \otimes u_{\mathcal{T}}, \quad \iff \quad \mathbf{U}_{i_0,\ldots,i_{\mathcal{T}}} = \prod_{t=0}^{\mathcal{T}} (u_t)_{i_t}, \tag{35}$$

where the vectors u_t *are given by* $u_t = \exp(G_t^T \lambda_t/\varepsilon)$ *for* $t = 0, \ldots, \mathcal{T}$*, with* $\lambda_t \in \mathbb{R}^{m_t}$ *denoting the Lagrange dual variable corresponding to the constraint* t *in* (34b)*. The optimal perturbation vectors may be expressed as* $\Delta_t = -\frac{1}{2\gamma_t}\lambda_t$*, for* $t = 0, 1, \ldots, \mathcal{T}$*. Furthermore, a Lagrange dual of the multi-marginal OMT problem* (34) *is given by*

$$\underset{\lambda_0,\ldots,\lambda_{\mathcal{T}}}{\text{maximize}} \quad -\varepsilon \langle \mathbf{K}, \mathbf{U} \rangle - \sum_{t=0}^{\mathcal{T}} \frac{1}{4\gamma_t} \|\lambda_t\|_2^2 + \sum_{t=0}^{\mathcal{T}} \lambda_t^T r_t. \tag{36}$$

where $\mathbf{U}$ *is given by* (35) *with* $u_t = \exp(G_t^T \lambda_t/\varepsilon)$ *for* $t = 0, \ldots, \mathcal{T}$*.*

Proof. See the appendix. □

With the result from Proposition 1, we are now ready to state the method for solving (34). Based on [31], and similarly to [21], we propose to solve the dual (36) by a block coordinate ascent.

Theorem 1. *Given an initial set of vectors* $\lambda_0, \ldots, \lambda_{\mathcal{T}}$*, iterate the following steps repeatedly for* $t \in \{0, \ldots, \mathcal{T}\}$*:*

- *Let*

$$v_t = P_t(\mathbf{K} \odot \mathbf{U})./u_t \tag{37}$$

  where $\mathbf{U} = u_0 \otimes u_1 \otimes \cdots \otimes u_{\mathcal{T}}$ *for the vectors* $u_t = \exp(G_t^T \lambda_t/\varepsilon)$*.*

- *Update the vector* λ_t *as the solution to*

$$G_t\left(v_t \odot \exp\left(\frac{G_t^T \lambda_t}{\varepsilon}\right)\right) + \frac{\lambda_t}{2\gamma_t} - r_t = 0. \tag{38}$$

Then, the vectors $\lambda_0, \ldots, \lambda_{\mathcal{T}}$ converge linearly to the unique optimal solution of the dual problem (36). *Furthermore, in the limit point of the iteration, the marginals of* $\mathbf{M}$ *are given directly as $\Phi_t = u_t \odot v_t$, for $t = 0, 1, \ldots, \mathcal{T}$.*

Proof of Theorem 1. We first show that the iterates in the theorem correspond to a block coordinate ascent in the dual problem (36). A block-coordinate ascent method is to iteratively update one of the dual variables as to maximize the objective function, while keeping the other variables fixed. Here, this is to iteratively update one λ_t, for $t = 0, 1, \ldots, \mathcal{T}$, such that it maximizes (36). The maximizing λ_t may be found as the root of the corresponding gradient. Therefore, note that substituting the explicit expression for the elements in $\mathbf{U}$, as detailed in Proposition 1, into the first term of the dual objective (36) yields

$$\sum_{i_0,\ldots,i_{\mathcal{T}}} \mathbf{K}_{i_0,\ldots,i_{\mathcal{T}}} \mathbf{U}_{i_0,\ldots,i_{\mathcal{T}}} = \sum_{i_0,\ldots,i_{\mathcal{T}}} \mathbf{K}_{i_0,\ldots,i_{\mathcal{T}}} \prod_{t=0}^{\mathcal{T}} \left(\exp(G_t^T \lambda_t / \varepsilon) \right)_{i_t}.$$

The gradient of this term with respect to λ_t can be written as

$$G_t(v_t \odot u_t),$$

with $u_t = \exp(G_t^T \lambda_t / \varepsilon)$ and the vector v_t defined as

$$\begin{aligned}
(v_t)_{i_t} &= \sum_{\substack{i_0,\ldots,i_{t-1},\\ i_{t+1},\ldots,i_{\mathcal{T}}}} K_{i_0,\ldots,i_{t-1},i_t,i_{t+1},\ldots,i_{\mathcal{T}}} \prod_{\substack{s=1 \\ s\neq t}}^{\mathcal{T}} (u_s)_{i_s} \\
&= \left(\sum_{\substack{i_0,\ldots,i_{t-1},\\ i_{t+1},\ldots,i_{\mathcal{T}}}} \mathbf{K}_{i_0,\ldots,i_{t-1},i_t,i_{t+1},\ldots,i_{\mathcal{T}}} \prod_{s=0}^{\mathcal{T}} (u_s)_{i_s} \right)_{i_t} / (u_t)_{i_t} \\
&= \left(P_t(\mathbf{K} \odot \mathbf{U}) \right)_{i_t} / (u_t)_{i_t}.
\end{aligned}$$

Hence, the gradient of the dual objective function (36) with respect to λ_t is

$$-G_t(v_t \odot u_t) - \lambda_t / (2\gamma_t) + r_t.$$

Furthermore, note that the dual objective function (36) is strongly concave. This can be seen by noting that the first term is concave, the second term is strongly concave, and the last term is linear. Since the dual is twice continuously differentiable, the block coordinate ascent method converges (locally) q-linearly to the unique global maximum [49, Theorem 2]. □

Remark 1. It is worth noting that, given the dual optimal variables λ_t, for $t = 0, 1, \ldots, \mathcal{T}$, obtained as the limit point of the iterations in Theorem 1, the primal optimal variables may be recovered using Proposition 1.

Remark 2. The number of constraints in (34) may be smaller than the number of modes in the mass transfer tensor $\mathbf{M}$. In cases where there are no constraints on the t-th projection of $\mathbf{M}$, the corresponding dual variable λ_t is set to zero. The method described in Theorem 1 may then be modified by setting $u_t = \mathbf{1}$ and only iterating over the smaller set of remaining vectors.

Recall that in the standard Sinkhorn iterations (31) and (32), the scaling factors u_t are iteratively updated as to satisfy the marginal constraints. The same property holds for the iterations in Theorem 1. Specifically, if λ_t is a solution to (38), then the implied transport plan in that iteration, i.e., $\mathbf{M} = \mathbf{K} \odot \mathbf{U}$, where $\mathbf{U} = u_0 \otimes u_1 \otimes \cdots \otimes u_{\mathcal{T}}$, with $u_t = \exp(G_t^T \lambda_t / \varepsilon)$, satisfies

$$G_t P_t(\mathbf{M}) = r_t + \Delta_t = r_t - \frac{1}{2\gamma_t}\lambda_t.$$

In the case of full information and exact matching in a marginal (i.e., $G_t = I$ and $\gamma_t = \infty$), the update (38) in Theorem 1 reduces to the Sinkhorn iterations (Algorithm 1), i.e.,

$$u_t \leftarrow r_t ./ v_t = u_t \odot r_t ./ P_t(\mathbf{K} \odot \mathbf{U}),$$

whereas in the case with only full information (i.e., $G_t = I$ and $0 < \gamma_t < \infty$), the update reduces to solving

$$v_t \odot \exp(\lambda_t / \varepsilon) + \lambda_t / (2\gamma_t) - r_t = 0. \tag{39}$$

This equation may be solved element-wise using the Wright omega function (see appendix B in [31]). This allows for computing the proximal operator of the regularized optimal mass transport [31], which is often used in first-order methods for non-smooth optimization. Similar expressions may also be obtained for other penalization terms than the squared ℓ_2-error, which relates to the entropic proximal operator [50]. Herein, we propose to solve (38) using Newton's method. Therefore, note that the corresponding Jacobian is given by

$$\frac{1}{\varepsilon} G_t \operatorname{diag}(v_t \odot u_t) G_t^T + \frac{1}{2\gamma_t} I.$$

Algorithm 2 Sinkhorn-Newton method for the multi-marginal optimal mass transport with partial information of the marginals.

Given: Initial guess λ_t, for $t = 0, \ldots, \mathcal{T}$; starting point t
$u_t \leftarrow \exp(G_t\lambda_t/\varepsilon)$ for $t = 0, \ldots, \mathcal{T}$
while Sinkhorn not converged **do**
  Construct v_t according to (37), i.e.,
  $v_t \leftarrow P_t(\mathbf{K} \odot \mathbf{U})./u_t$
  while Newton not converged **do**
    $f \leftarrow G_t \operatorname{diag}(v_t)u_t - r_t + 1/(2\gamma_t)\lambda_t$
    $df \leftarrow (1/\varepsilon)G_t \operatorname{diag}(v_t \odot u_t)G_t^T + 1/(2\gamma_t)I$
    $\Delta\lambda \leftarrow -df \backslash f$
    $\lambda_t \leftarrow \lambda_t + \eta\Delta\lambda$, with η determined by a linesearch
    $u_t \leftarrow \exp(G_t\lambda_t/\varepsilon)$
  end while
  $t \leftarrow t + 1 \pmod{\mathcal{T}+1}$
end while
return $\mathbf{M} = \mathbf{K} \odot \mathbf{U}$

The full method for solving (34) is summarized in Algorithm 2. It may be noted that the multi-marginal transport plan $\mathbf{M}$ may for some problems be too large to be stored in the memory. However, when used in modeling applications, one is primarily interested in being able to study projections of $\mathbf{M}$, which, due to the structure of the transport plan, do not require $\mathbf{M}$ to be constructed explicitly.

Remark 3. Sufficiently close to the optimal solution, the quadratic approximation of the dual objective underlying the Newton method becomes increasingly accurate. The inner Newton method for solving (38) then converges in the first iteration. In the authors' experience, this is typically achieved within the first few outer Sinkhorn iterations.

The computational bottleneck of Algorithm 2 is the construction of the vectors v_t in (37), requiring the computation of $P_t(\mathbf{K} \odot \mathbf{U})$. Structures in the cost tensor $\mathcal{T}$ may allow for efficiently computing these projections, as we will see in Section 5. With the results from Section 5, Theorem 1 provides an efficient scheme for implementing the spatial spectral estimators in Sections 3.1 - 3.3.

5 Efficient computations for structured multi-marginal OMT problems

In many applications, the size of the mass transport tensor $\mathbf{M}$, and thus of the auxiliary tensors $\mathbf{K}$ and $\mathbf{U}$, may be too large to manipulate directly. It is thus crucial to utilize additional structures in the problem whenever this is possible. In this section, some common examples for structures in the cost tensor $\mathbf{C}$ are described, allowing for efficient computation of subproblems in the multi-marginal Sinkhorn algorithm (Algorithm 1). The computational bottleneck of the Sinkhorn scheme, as well as the method for partial information of the marginals (Algorithm 2), is the computation of $P_t(\mathbf{K} \odot \mathbf{U})$, where both the elementwise multiplication $\mathbf{K} \odot \mathbf{U}$ and the application of the projection operator P_t may be expensive. However, in several cases of interest, the choice of cost function, c, induces structure in the tensor $\mathbf{K}$ that may be exploited in order to dramatically increase the efficiency of computing $P_t(\mathbf{K} \odot \mathbf{U})$. Such structures appears in e.g., barycenter problems [19], tracking problems [20, 32], and certain hidden Markov chain formulations [51]. Here, we will consider cases relevant for the applications presented in Section 3. In fact, for these cases, the computation of $P_t(\mathbf{K} \odot \mathbf{U})$ may be performed by sequences of matrix-vector multiplications as detailed in the propositions in the following subsections.

5.1 Sequentially decoupling cost function

Consider the multi-marginal OMT problem on the form (15) where the cost function $\mathcal{C} : \mathcal{X} \times \mathcal{X} \to \mathbb{R}_+$ decouples sequentially according to (22), which may be utilized in order to model sequential tracking over time. The following proposition shows how this special case allows for efficient computations of the projections $P_t(\mathbf{K} \odot \mathbf{U})$ and $P_{t_1,t_2}(\mathbf{K} \odot \mathbf{U})$ in the discrete OMT problem.

Proposition 2. *Let the elements of the cost tensor* $\mathbf{C}$ *in* (28) *be of the form*

$$\mathbf{C}_{i_0,\dots,i_\mathcal{T}} = \sum_{t=1}^{\mathcal{T}} C_{i_{t-1},i_t},$$

for a cost matrix $C \in \mathbb{R}^{n\times n}$*, and let* $\mathbf{K} = \exp(-\mathbf{C}/\varepsilon)$*,* $\mathbf{U} = (u_0 \otimes u_1 \otimes \cdots \otimes u_\mathcal{T})$*.*

Then, for $0 \leq t \leq \mathcal{T}$,

$$P_t(\mathbf{K} \odot \mathbf{U}) = \left(u_0^T K \operatorname{diag}(u_1) K \dots K \operatorname{diag}(u_{t-1}) K\right)^T \odot u_t \\ \odot \left(K \operatorname{diag}(u_{t+1}) K \dots K \operatorname{diag}(u_{\mathcal{T}-1}) K u_{\mathcal{T}}\right),$$

where $K = \exp(-C/\varepsilon)$, *and, for* $0 \leq t_1 < t_2 \leq \mathcal{T}$,

$$P_{t_1,t_2}(\mathbf{K} \odot \mathbf{U}) = \operatorname{diag}\left(u_0^T K \operatorname{diag}(u_1) K \cdots K \operatorname{diag}(u_{t_1-1}) K\right) \\ \odot \operatorname{diag}(u_{t_1}) K \operatorname{diag}(u_{t_1+1}) K \dots K \operatorname{diag}(u_{t_2}) \\ \odot \operatorname{diag}\left(K \operatorname{diag}(u_{t_2+1}) K \dots K \operatorname{diag}(u_{\mathcal{T}-1}) K u_{\mathcal{T}}\right). \tag{40}$$

Proof. See the appendix. □

It is worth noting that the computation of $P_t(\mathbf{K} \odot \mathbf{U})$ requires only $\mathcal{T}$ matrix-vector multiplications of the form Ku_τ, i.e., $\mathcal{O}(\mathcal{T} n^2)$ operations, whereas the complexity for a brute-force approach is $\mathcal{O}(n^{\mathcal{T}+1})$ operations. The same holds for the multiplication of $P_{t_1,t_2}(\mathbf{K} \odot \mathbf{U})$ with a vector. If several projections are computed iteratively, as, e.g., in Algorithm 1 or Algorithm 2, intermediate results of the factors to the left and right of u_t, for $t = 0, \dots, \mathcal{T}$, may be stored. Then each iteration requires only one matrix vector multiplication, that is $\mathcal{O}(n^2)$ operations. Furthermore, if the domains of the marginals have cost functions that decouple then the efficiency of the matrix vector computations can be improved even further, as described in the following remark.

Remark 4. In many multidimensional problems, the domain X can be represented as a direct product $X = X_1 \times X_2 \times \cdots \times X_N$, with $|X_i| = n_i$. If, in addition, the cost decouples in the same way, i.e., for two points $\mathrm{x}_0 = \left(\mathrm{x}_0^{(1)}, \mathrm{x}_0^{(2)}, \dots, \mathrm{x}_0^{(N)}\right) \in X$ and $\mathrm{x}_1 = \left(\mathrm{x}_1^{(1)}, \mathrm{x}_1^{(2)}, \dots, \mathrm{x}_1^{(N)}\right) \in X$,

$$c(\mathrm{x}_0, \mathrm{x}_1) = \sum_{i=1}^{N} c_i\left(\mathrm{x}_0^{(i)}, \mathrm{x}_1^{(i)}\right),$$

where $c_i : X_i \times X_i \to \mathbb{R}$ is a cost function in the i-th dimension, for $i = 1, \dots, N$, then the matrix $K = \exp(-C/\varepsilon) \in \mathbb{R}^{n \times n}$, where $n = \prod_{i=1}^{N} n_i$, can be decoupled as a tensor product $K = K_1 \otimes K_2 \otimes \cdots \otimes K_N$. Multiplications with K can then be done one dimension at a time. Practically, this can be efficiently implemented

using specialized tensor toolboxes, e.g., with the command `tmprod` of the Matlab toolbox Tensorlab [52]. Thus, the multiplication Ku requires only $\mathcal{O}(n\sum_i n_i)$ operations instead of $\mathcal{O}(n^2)$ which is required when using standard multiplication. From a memory perspective, this is also important since K does not need to be computed and stored. In particular, this holds for a regular rectangular grid in $\mathbb{R}^N$ with cost function $c(x,y) = \|x-y\|_p^p$, for some $p \geq 1$, which is the typical setup for the Wasserstein metric.

5.2 Centrally decoupling cost function

In order to solve OMT barycenter formulations, we may use centrally decoupled cost functions $\boldsymbol{C}$ according to (19) to model transport between a common barycenter and several marginal distributions. The following proposition describes the projection computations for this case.

Proposition 3. *Let the elements of the cost tensor* $\mathbf{C}$ *in* (28) *be of the form*

$$\mathbf{C}_{i_0,\ldots,i_J} = \sum_{j=1}^{J} C_{i_0,i_j},$$

for a cost matrix $C \in \mathbb{R}^{n\times n}$*, and let* $\mathbf{K} = \exp(-\mathbf{C}/\varepsilon)$*,* $\mathbf{U} = (u_0 \otimes u_1 \otimes \cdots \otimes u_T)$*. Then,*

$$P_0(\mathbf{K}\odot\mathbf{U}) = u_0 \odot \bigodot_{\ell=1}^{J} (Ku_\ell), \tag{41}$$

$$P_j(\mathbf{K}\odot\mathbf{U}) = u_j \odot K^T \left(u_0 \odot \bigodot_{\ell=1,\ell\neq j}^{J} (Ku_\ell) \right), \tag{42}$$

for $j = 1,\ldots,J$*, where* $K = \exp(-C/\varepsilon)$*. For* $j, j_1, j_2 = 1,\ldots,J$*, and* $j_1 \neq j_2$*, the pair-wise projections are given by*

$$P_{0,j}(\mathbf{K}\odot\mathbf{U}) = \operatorname{diag}\left(u_0 \odot \bigodot_{\substack{\ell=1\\ \ell\neq j}}^{J} (Ku_\ell) \right) K \operatorname{diag}(u_j), \tag{43}$$

$$P_{j_1,j_2}(\mathbf{K}\odot\mathbf{U}) = \operatorname{diag}(u_{j_1}) K^T \operatorname{diag}\left(p_{j_1,j_2}\right) K \operatorname{diag}(u_{j_2}), \tag{44}$$

where

$$p_{j_1,j_2} = u_0 \odot \bigodot_{\substack{\ell=1 \\ \ell \neq j_1,j_2}}^{J} (Ku_\ell).$$

Proof. See the appendix. □

Similarly to the setting of a sequentially decoupling cost, the computation of a projection $P_j(\mathbf{K} \odot \mathbf{U})$, for any $j = 0, \ldots, J$, in the case of a centrally decoupling cost, involves J matrix vector multiplications. It thus requires $\mathcal{O}(Jn^2)$ operations, as opposed to $\mathcal{O}(n^{J+1})$ in the brute force approach. For the iterative schemes in Algorithm 1 and Algorithm 2, storing intermediate results decreases the complexity in each iteration to at most two matrix vector multiplications, i.e., $\mathcal{O}(n^2)$ operations.

5.3 Sequential and central decoupling

Consider a cost function $\boldsymbol{\mathcal{C}}$ being a combination of sequentially and centrally decoupled costs, according to (25), which in Section 3.3 was applied to modeling tracking of barycenters over time. For this case, the following proposition holds.

Proposition 4. *Let the elements of the cost tensor* $\mathbf{C}$ *in* (28) *be of the form*

$$\mathbf{C}_{(i_{(t,j)} \mid (t,j) \in \Lambda)} = \sum_{t=1}^{\mathcal{T}} C_{i_{(t-1,0)}, i_{(t,0)}} + \sum_{t=1}^{\mathcal{T}} \sum_{j=1}^{J} \tilde{C}_{i_{(t,0)}, i_{(t,j)}},$$

for cost matrices $C \in \mathbb{R}^{n \times n}$ *and* $\tilde{C} \in \mathbb{R}^{n \times \tilde{n}}$*, and define* $\Lambda = \{(t,j) \mid t \in \{0,1,\ldots,\mathcal{T}\}, j \in \{0,1 \ \ldots, J\}\}$*. Furthermore, let* $\mathbf{K} = \exp(-\mathbf{C}/\varepsilon)$*, and let* $\mathbf{U} = \bigotimes_{(t,j)\in\Lambda} u_{(t,j)}$*. Then, for the central marginals, corresponding to* (t,j) *such that* $j = 0$*,*

$$\begin{aligned} P_{(t,0)}(\mathbf{K} \odot \mathbf{U}) = &\left(p_0^T K \operatorname{diag}(p_1) K \ldots K \operatorname{diag}(p_{t-1}) K\right)^T \odot p_t \\ &\odot \left(K \operatorname{diag}(p_{t+1}) K \ldots K \operatorname{diag}(p_{\mathcal{T}-1}) K p_{\mathcal{T}}\right), \end{aligned} \tag{45}$$

where $p_t = u_{(t,0)} \odot \bigodot_{j=1}^{J} \tilde{K} u_{(t,j)}$ *and* $K = \exp(-C/\varepsilon)$*,* $\tilde{K} = \exp(-\tilde{C}/\varepsilon)$*. Furthermore, for the non-central marginals, i.e.,* (t,j) *with* $j = 1, 2, \ldots, J$*, it holds*

that

$$P_{(t,j)}(\mathbf{K} \odot \mathbf{U}) = u_{(t,j)} \odot \tilde{K}^T \Big((p_0^T K \operatorname{diag}(p_1) K \dots K \operatorname{diag}(p_{t-1}) K)^T \\ \odot (p_t ./ (\tilde{K} u_{(t,j)})) \\ \odot (K \operatorname{diag}(p_{t+1}) K \dots K \operatorname{diag}(p_{\mathcal{T}-1}) K p_{\mathcal{T}}) \Big). \tag{46}$$

Proof. See the appendix. □

Computing one vector p_t, for any $t = 0, \dots, \mathcal{T}$, requires J matrix vector multiplications, i.e., $\mathcal{O}(Jn\tilde{n})$ operations. In order to compute the projection $P_{(t,0)}(\mathbf{K} \odot \mathbf{U})$, a total of $\mathcal{T}+1$ vectors of this form have to be computed, as well as $\mathcal{T}$ aditional matrix vector products of complexity $\mathcal{O}(n^2)$ each. This results in a total of $\mathcal{O}(\mathcal{T}Jn\tilde{n} + \mathcal{T}n^2)$ operations. The complexity is the same for the other projections, $P_{(t,j)}(\mathbf{K} \odot \mathbf{U})$, for $j = 1, \dots, J$, and significantly lower than the direct computation of the projections, which requires $\mathcal{O}(n^{(J+1)(\mathcal{T}+1)})$ operations. It is worth noting that in an iterative scheme such as Algorithm 1 or Algorithm 2, when cycling through $j = 0, 1, \dots, J$, for a given t, one does not have to re-compute p_τ for $\tau \neq t$, i.e., the multiplicative factors to the left and right of $p_t ./ (K u_{(t,j)})$ in (46) have to be updated only once for each t. Furthermore, as noted in Section 5.1, these factors may be stored so that updating them for a given t requires only one multiplication with K. Moreover, as noted in Section 5.2, the updates of p_t while cycling through $j = 0, \dots, J$ involve at most two multiplications with $\tilde{K}$ in each iteration. Thus, computing the projections $P_{(t,j)}(\mathbf{K} \odot \mathbf{U})$, for all $j = 0, \dots, J$ and a fixed t, requires $\mathcal{O}(n^2 + Jn\tilde{n})$ operations.

Remark 5. The pairwise projections for index pairs $(t_1, 0)$ and $(t_2, 0)$, i.e., indices corresponding to the central marginals, may be computed analogously to Proposition 2. That is,

$$P_{(t_1,0),(t_2,0)}(\mathbf{K} \odot \mathbf{U}) = \operatorname{diag}\Big(p_0^T K \operatorname{diag}(p_1) K \dots K \operatorname{diag}(p_{t_1-1}) K \Big) \\ \odot \operatorname{diag}(p_{t_1}) K \operatorname{diag}(p_{t_1+1}) K \cdots K \operatorname{diag}(p_{t_2}) \\ \odot \operatorname{diag}\big(K \operatorname{diag}(p_{t_2+1}) K \dots K \operatorname{diag}(p_{\mathcal{T}-1}) K p_{\mathcal{T}} \big).$$

In case the index pairs are (t, j_1) and (t, j_2), i.e., when the indices correspond to the same central marginal, the pairwise projections can be computed similarly to

Proposition 3. Specifically, the bi-marginal projection for a central marginal and a non-central marginal corresponding to the same time t is given by

$$P_{(t,0),(t,j)}(\mathbf{K} \odot \mathbf{U}) = \operatorname{diag}(u_{(t,0)}) \operatorname{diag}\left(\tilde{K}^T \rho_{(t,0),(t,j)}\right) \tilde{K} \operatorname{diag}(u_{(t,j)}),$$

where

$$\begin{aligned}\rho_{(t,0),(t,j)} &= (p_0^T K \operatorname{diag}(p_1) K \ldots K \operatorname{diag}(p_{t-1}) K)^T \\ &\odot (p_t ./ (u_{(t,0)} \odot \tilde{K} u_{(t,j)})) \odot (K \operatorname{diag}(p_{t+1}) K \ldots K \operatorname{diag}(p_{\mathcal{T}-1}) K p_{\mathcal{T}}),\end{aligned}$$

whereas the bi-marginal projection for two non-central marginals corresponding to the same time t may be expressed as

$$P_{(t,j_1),(t,j_2)}(\mathbf{K} \odot \mathbf{U}) = \operatorname{diag}(u_{(t,j_1)}) \tilde{K}^T \operatorname{diag}(\rho_{(t,j_1),(t,j_2)}) \tilde{K} \operatorname{diag}(u_{(t,j_2)}),$$

where

$$\begin{aligned}\rho_{(t,j_1),(t,j_2)} &= (p_0^T K \operatorname{diag}(p_1) K \ldots \operatorname{diag}(p_{t-1}) K)^T \\ &\odot (p_t ./ (\tilde{K} u_{(t,j_1)} \odot \tilde{K} u_{(t,j_2)})) \odot (K \operatorname{diag}(p_{t+1}) \ldots K \operatorname{diag}(p_{\mathcal{T}-1}) K p_{\mathcal{T}}).\end{aligned}$$

For arbitrary index pairs (t, j), expressions for the pairwise projections may be derived in a similar way.

6 Numerical results

In this section, we demonstrate the behavior of the proposed multi-marginal OMT problem in (16) when applied to the discussed spatial spectral estimation problems. Throughout, the problem in (16) is approximated by the discrete counterpart in (34), and solved using Algorithm 2. For the implementation of the projection operators, P_t, we use the results from the propositions in Section 5. For the example in Section 6.4, we provide a dimensionality analysis of the corresponding OMT problem, illustrating the benefit of the presented computational tools.

6.1 Tracking with static and dynamical models

In order to illustrate the different properties of the static and dynamic OMT formulations, we consider a DoA tracking example with two moving signal sources.

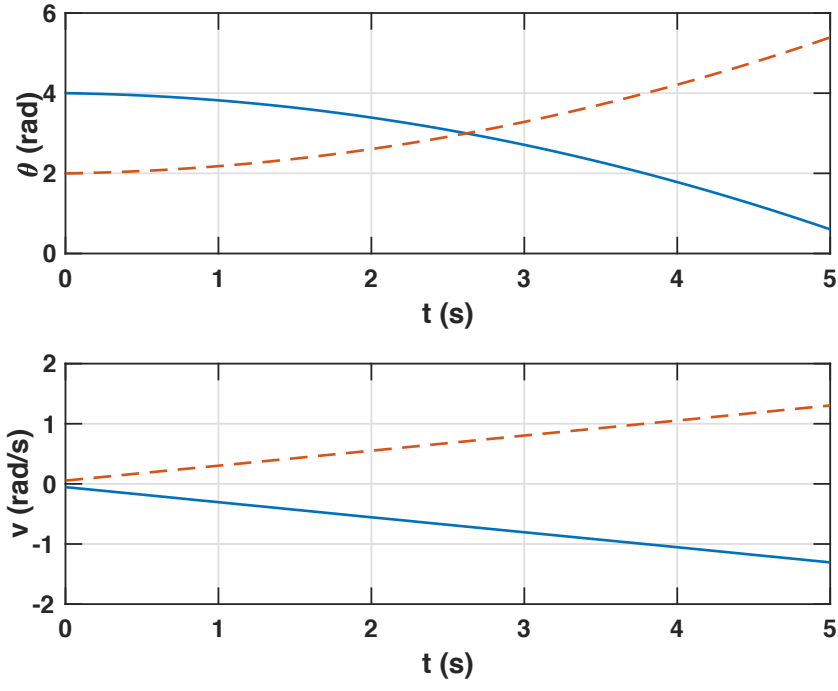

Figure 1: Ground truth of the DoA tracking example with two moving targets. Top panel: the target DoAs, in radians, as a function of time. Bottom panel: the target velocities, in radians per second, as a function of time.

Specifically, we consider a ULA consisting of 5 sensors with half-wavelength spacing, measuring the superposition of the source signals, modeled as independent Gaussian processes, together with a spatially white complex Gaussian sensor noise. The signal to noise ratio (SNR), defined as

$$\text{SNR} = 10 \log_{10}(\sigma^2_{signal}/\sigma^2_{noise}),$$

where σ^2_{signal} and σ^2_{noise} are the signal and noise powers, respectively, is 10 dB. The trajectories of the two sources are displayed in the top panel of Figure 1, with the bottom panel showing the velocities. At six different time instances, evenly spaced in time throughout the trajectories, we collect 25 array signal snapshots, from which the array covariance is estimated using the sample covariance matrix. Using the resulting sequence of covariance matrices, we then attempt to reconstruct the

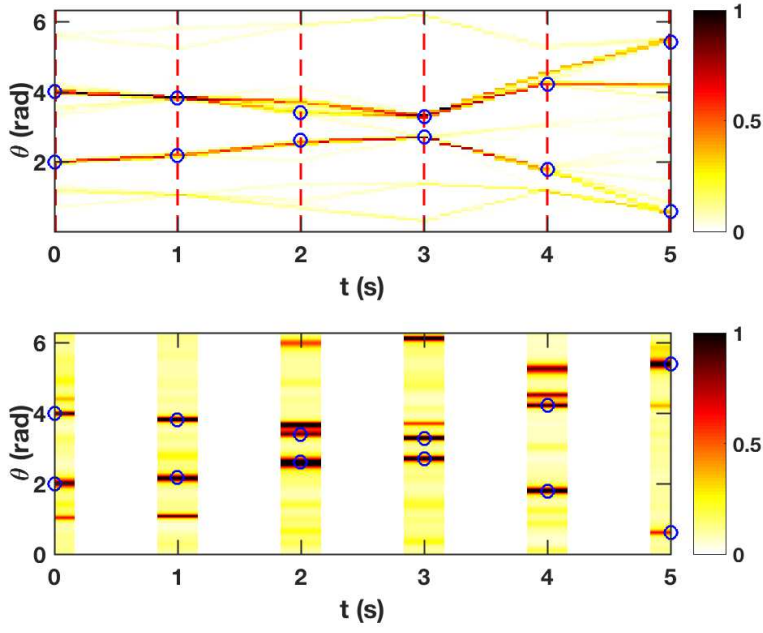

Figure 2: Top panel: reconstructed DoA spectrum using a static cost function in the OMT formulation, based on six observations of the array covariance matrix. The observation times are marked in dashed lines. The ground truth trajectories at the observation times are marked by rings. Bottom panel: estimates obtained using the Capon estimator applied to the individual covariance matrices.

target trajectories using the static and dynamic OMT problems. For the static OMT problem, we use the cost function $c(\vartheta, \varphi) = |\vartheta - \varphi|^2$, for $\vartheta, \varphi \in (-\pi, \pi]$. For the dynamical model, we introduce a latent velocity state, v, such that $\dot{\vartheta}(t) = v(t)$, and use the state space representation in (12), where

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 0 & 1 \end{bmatrix}^T, \quad F = \begin{bmatrix} 1 & 0 \end{bmatrix},$$

with the state vector being formed as $x(t) = \begin{bmatrix} \vartheta(t) & v(t) \end{bmatrix}^T$, i.e., the same model as in Example 1. Note here that the F matrix reflects the fact that only the angle, ϑ, is directly observable in the array covariance matrix. The choice of B implies

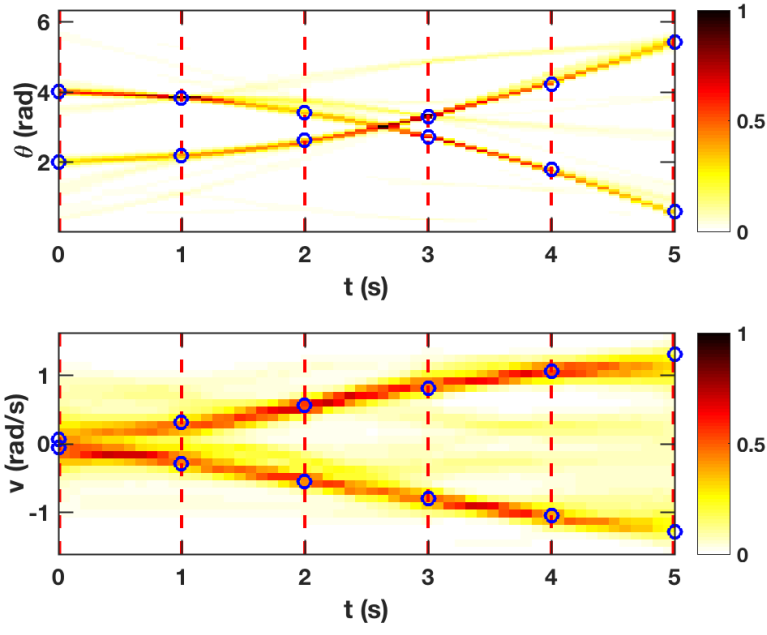

Figure 3: Reconstructed DoA and velocity spectra using a dynamic cost function in the OMT formulation, based on six observations of the array covariance matrix. The observation times are marked in dashed lines. The ground truth trajectories at the observation times are marked by rings. Top panel: reconstructed DoA spectrum. Bottom panel: reconstructed velocity spectrum.

that the angle may only be influenced via the velocity, i.e., through acceleration. Thus, the resulting dynamic cost function, as defined in (8), penalizes transport requiring acceleration, whereas the cost function for the static problem in contrast penalizes transport requiring velocity. In the discrete implementations of the methods, we use 100 grid points to represent the angle, ϑ, and the dynamical model uses 30 grid points for the velocity, v. Also, we use the regularization parameters $\varepsilon = 10^{-1}$, and common parameters $\gamma = 30$ and $\gamma = 5$ for all marginals for the static and dynamic OMT models, respectively. The results for the static and dynamical models are displayed in the top panels of Figures 2 and 3, respectively. Here, the trajectories in between the observation times, which are

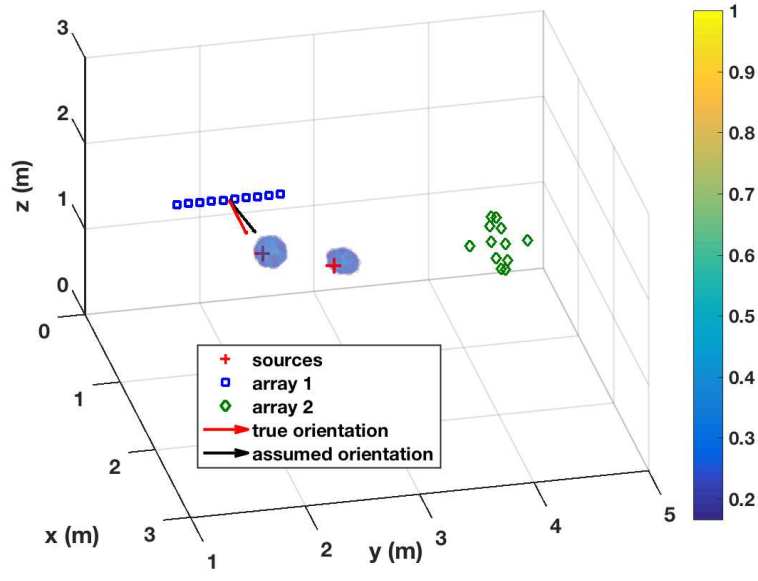

Figure 4: Scenario for localization in 3D using two sensor arrays and two signal sources. The true and assumed orientations of the first array are indicated by arrows. The three dimensional spectral estimate obtained using the formulation in (20) is superimposed.

indicated by vertical dashed lines, are reconstructed using the interpolation procedure presented in Section 2.2, as defined by the obtained multi-marginal transport plan. The required bi-marginal transport plans describing the mass transfer between the respective margins, i.e., consecutive observation points, are computed as the bi-marginal projections detailed in Proposition 2. To simplify comparison, the ground truth trajectories from Figure 1, at the observation times, are superimposed in the plots. Also, the bottom panel of Figure 2 displays the spectral estimates obtained by applying the Capon estimator [53] to the individual covariance matrix estimates. As may be noted, the Capon estimate contains several spurious peaks, caused by the noisy measurements. In contrast, as seen in Figure 2, the static model is able to produce reasonable spectral estimates for the observation times. However, the reconstructed trajectories fail to model the crossing of the paths of the targets. This is not unexpected: as the static OMT model

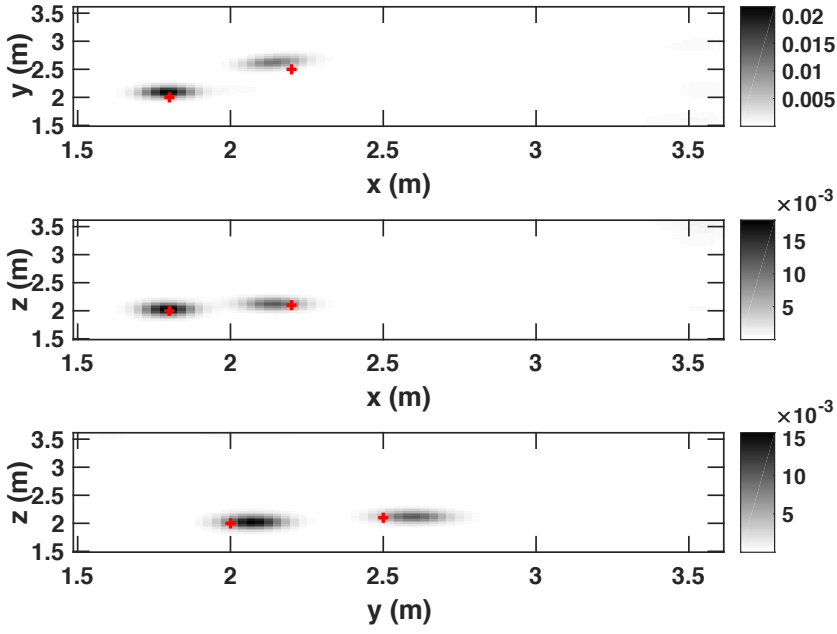

Figure 5: Two-dimensional projections of the three-dimensional spectral estimate obtained using the formulation in (20) for the estimation scenario in Figure 4.

penalizes movement, i.e., velocity, the cost of transport is smaller if the targets instead change course (note the trajectory between $t = 2$ and $t = 3$). In contrast, the dynamical formulation, in which movement is expected, is able to produce considerably more accurate estimates, as may be seen in Figure 3, including the crossing of the paths of the targets. Note also that the dynamical formulation is able to reconstruct the spectrum also for the hidden velocity state.

6.2 Sensor fusion in 3-D - audio example

In order to illustrate the applicability of the multi-marginal OMT formulation in (16), in the form of (20), for performing sensor fusion, we consider a three-dimensional (3-D) localization problem, in which two sensor arrays observe two signal sources. The sources are modeled as localized speech sources, with the source signals being taken from babble noise excerpts from [54]. The arrays consist of ten sensors each, arranged as a ULA for the first array, and as points on a

circle, with two additional sensors perpendicular to the plane of the circle, for the second array. The sensor spacing for the ULA is 0.1 meters, and the radius for the circular array is 0.25 meters. The scenario is shown in Figure 4. The measured sensor signals are generated using acoustic impulse responses obtained from the randomized image method [55], using the room dimensions $4.3 \times 6.9 \times 2.6$ meters. As the sources are broadband, we consider processing in the short-time Fourier transform (STFT) domain, i.e., equivalent to narrow-band filtering of the signal. Specifically, the signal is sampled at 16 kHz, and the STFT representation consists of 256 frequency bins, where each frame is constructed using 16 ms of the signal, using a Hann window with 50% overlap. We then compute estimates of the covariance matrices for the respective arrays corresponding to the frequency 2437.5 Hz, i.e., for the wavelength 0.1395 meters, using the sample covariance estimate constructed from 100 signal snapshots. The proposed barycenter method in (20) is then used to form a joint spatial spectral estimate using the two estimated covariance matrices, with the cost function $c(x_0, x_1) = \|x_0 - x_1\|_2^2$, for $x_0, x_1 \in \mathbb{R}^3$. For the discrete implementation in (34), we use a uniform gridding of the cube $[1.5, 3.5]^3$, using $n = 75$ points in each dimension. The regularization parameters are $\gamma = 1$, common for all marginals, and $\varepsilon = 10^{-2}$. Also, in order to illustrate the geometrical properties of the proposed formulation, we assume that the geometry of the ULA is only approximately known. Specifically, the assumed array geometry corresponds to a rotation in the $x - y$ plane of the true array. The obtained spectral estimate is superimposed in Figure 4. As can be seen, the obtained estimate implies a spatial translation of the actual sources. However, it may be noted that the obtained estimate clearly identifies two targets, i.e., the error in the array orientation only results in a spatial perturbation, but no artifacts in the form of, e.g., spurious sources. This is also illustrated in Figure 5, showing the projection of the three-dimensional spectrum onto the two dimensional subspaces. As can be seen from the second and third panel, the position in the z-coordinate is unbiased, as the rotation only takes place in the $x - y$ plane.

6.3 Sensor fusion 2-D - simulation study

As illustrated above, the proposed multi-marginal formulation is able to produce easily interpretable spectral estimates when used in 3-D sensor fusion scenarios, despite having erroneous information of the array geometry. Elaborating on this, we here conduct a Monte Carlo simulation study as to investigate the behavior of the spectral barycenter as a function of the array alignment error, and com-

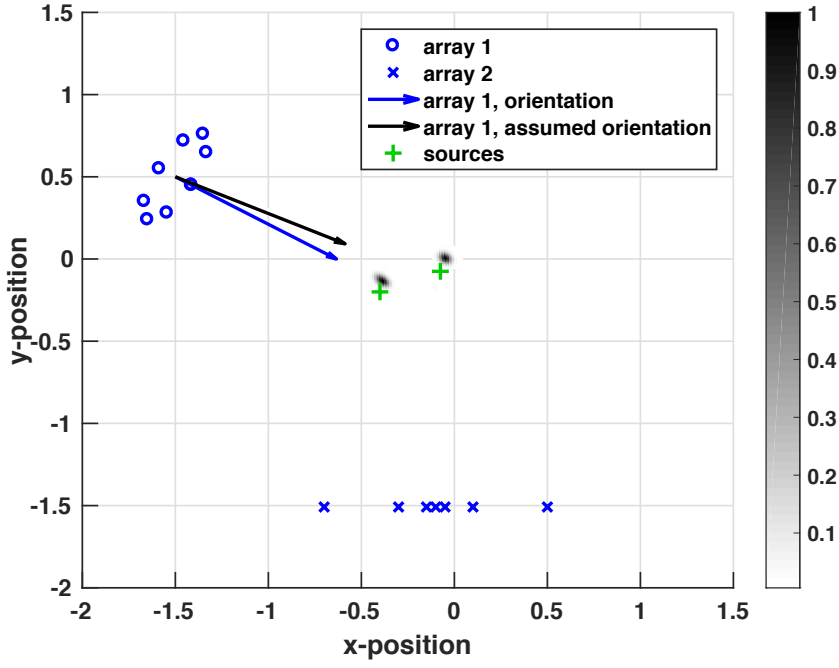

Figure 6: Spectral estimate as given by the multi-marginal barycenter formulation in (20). The alignment error is 6.7 degrees.

pare to other spectral estimation methods applicable to scenarios in which only the covariance matrices of individual sensor arrays, but not the inter-array covariances, are available. Specifically, consider a 2-D localization scenario, in which two uncorrelated sources impinge on two sensor arrays; one ellipsoidal shaped array consisting of 8 sensors, and one linear array consisting of 7 sensors. The wavefronts are here modeled as being circular. The scenario is shown in Figure 6. As may be seen, similar to the three-dimensional example, an unknown rotation is introduced to the ellipsoidal array, here varying this rotation between 0 and 10 degrees. In the Monte Carlo simulation, 100 realizations are generated for each considered rotation angle. In each realization, the locations of the signal sources are randomized uniformly on the square $[-0.5, 0.5] \times [-0.5, 0.5]$. The sources are modeled as uncorrelated circularly symmetric Gaussian white noises with variance 100, and a spatially white noise with variance 1 is added to the sensor measurements. The wavelength of the impinging waves is twice that of the smallest

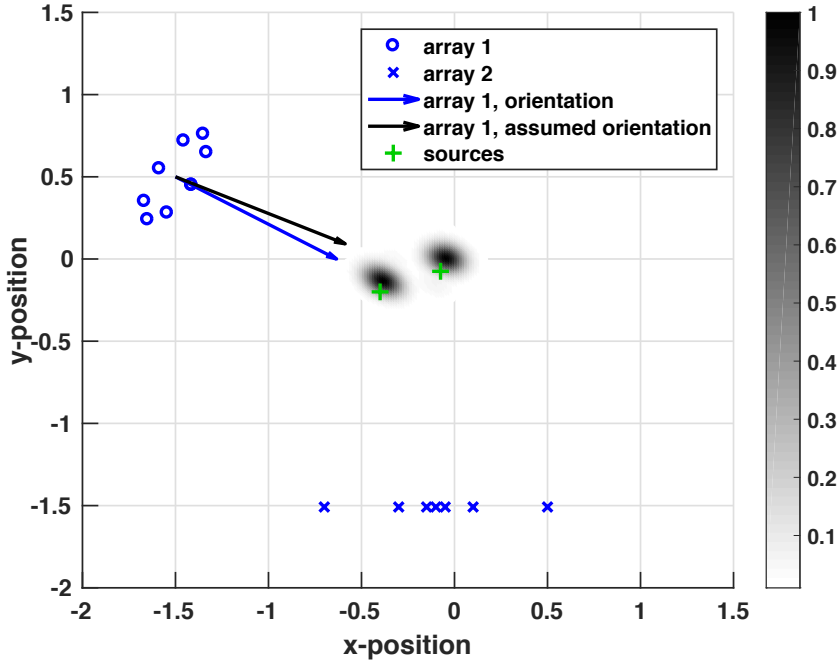

Figure 7: Spectral estimate as given by the pair-wise barycenter formulation in (18). The alignment error is 6.7 degrees.

sensor spacing in the linear array. We here consider spectral estimates obtained by the multi-marginal barycenter formulation in (20), as approximated by the discretization in (34). We also consider the formulation in (18) which is discretized and where the entropy regularization is employed separately for each pair-wise transport plan (see [21] for details on the implementation of this problem). For both methods, we utilize the cost function $c(\mathrm{x}_k, \mathrm{x}_\ell) = \|\mathrm{x}_k - \mathrm{x}_\ell\|_2^2$, for grid points $\mathrm{x}_k, \mathrm{x}_\ell \in \mathbb{R}^2$. We grid the square $[-1, 1]^2$ uniformly using $n = 100$ points in each dimension, and the algorithm parameters are $\gamma = 0.01$, common for all marginals, and $\varepsilon = 10^{-3}$ and $\varepsilon = 5 \cdot 10^{-3}$ for the multi-marginal and pair-wise formulations, respectively. The array covariance matrices are estimated using the sample covariance matrix from 500 signal snapshots. As comparison, we also apply the non-coherent MUSIC and MVDR estimators, as described in [56], as well as the least-squares (LS) estimator from [57], and the non-coherent SPICE estimator from [58] to the estimated covariance matrices in order to obtain estimates of

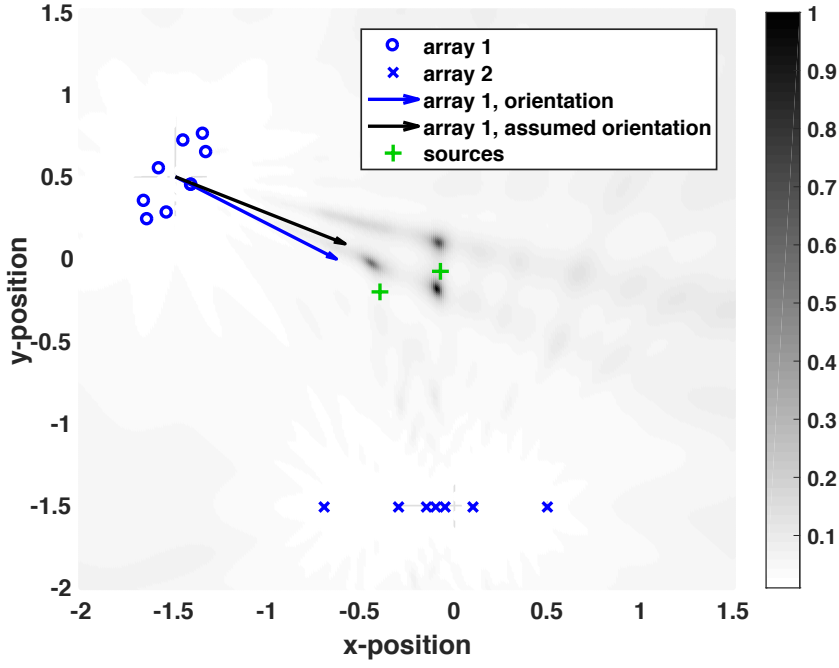

Figure 8: Pseudo-spectrum as given by non-coherent MUSIC. The alignment error is 6.7 degrees.

the spatial spectrum. Also included are estimates produced by the robust estimators RARE [59], extended to near-field localization [60], as well as the method introduced in [61], here referred to as AW19. It should here be noted that these robust estimators rely on knowledge of the joint covariance matrix of the sensor arrays, and have therefore, in contrast to the proposed method as well as the other comparison methods, been provided with the inter-array cross-covariances.

Figures 6, 7, and 8 provide illustrations of the behavior of the considered methods for the case of 6.7 degrees misalignment, for the multi-marginal formulation, the pair-wise regularized barycenter, and the MUSIC estimator, respectively. Comparing the two barycenter formulations, it may be noted that the estimate provided by the multi-marginal representation is considerably more concentrated, as compared to the formulation with pair-wise regularization. This is due to the optimization problem in (34) being more well-conditioned, allowing for smaller entropy regularization while still retaining numerical stability. It should how-

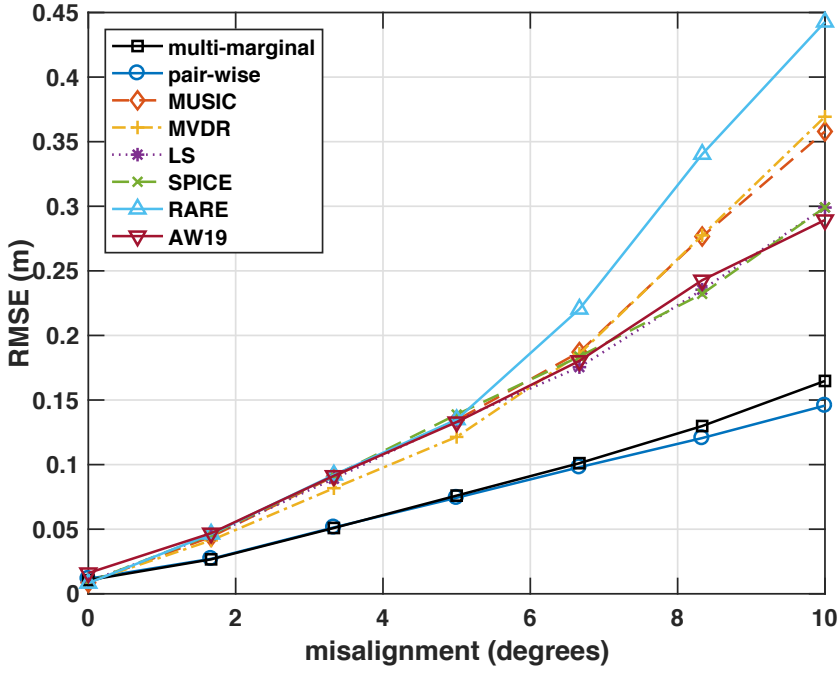

Figure 9: Error in location of spectral peaks, as function of the misalignment angle.

ever be noted that the values of the parameter ε does not have exactly the same meaning for the two formulations, i.e., the problems are not equivalent even for identical parameter values. It may also be noted that the spectral estimates obtained using both barycenter formulations only imply a spatial perturbation of the signal sources. In contrast, the pseudo-spectral estimates obtained using the MUSIC estimator contain spurious peaks, in addition to larger deviations from the true source locations.

The results from the full simulation study are shown in Figure 9, displaying the root mean squared error (RMSE) for the deviation of the spectral peaks to the true source locations. It may here be noted that for the case of no misalignment, the barycenter formulations produce estimates deviating slightly more from the ground truth than the comparison methods. However, as the misalignment increases, the RMSE for the barycenter estimators increase more slowly, indicating a greater robustness. This is not unexpected, as the formulations in (20) and (18)

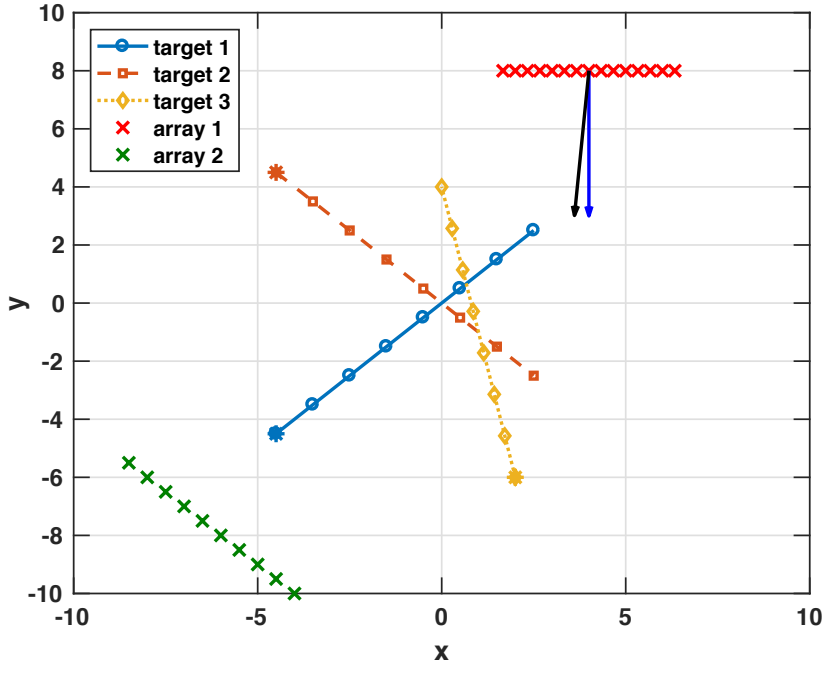

Figure 10: Scenario for the barycenter tracking problem with three moving sources. The start of the trajectory for each source is indicated by an asterisk (∗). Note here that one of the sensor arrays is slightly rotated, as indicated by the orientation arrows.

allows for perturbations on the underlying domain, whereas the optimization criteria for the LS and SPICE estimators are related to L_2-distances for the spectra. It may be noted that the robust RARE and AW19 estimators perform similarly to the other comparison methods, arguably due to the quite substantial calibration errors introduced by the array rotations. The sensitivity of RARE to the presence of array rotations and not only pure translations, although for the far-field case, has earlier been noted in [62].

6.4 Tracking of barycenters

As an illustration of the barycenter tracking formulation in Section 3.3, we consider a 2-D localization scenario in which two ULAs, each consisting of 15 sensors,

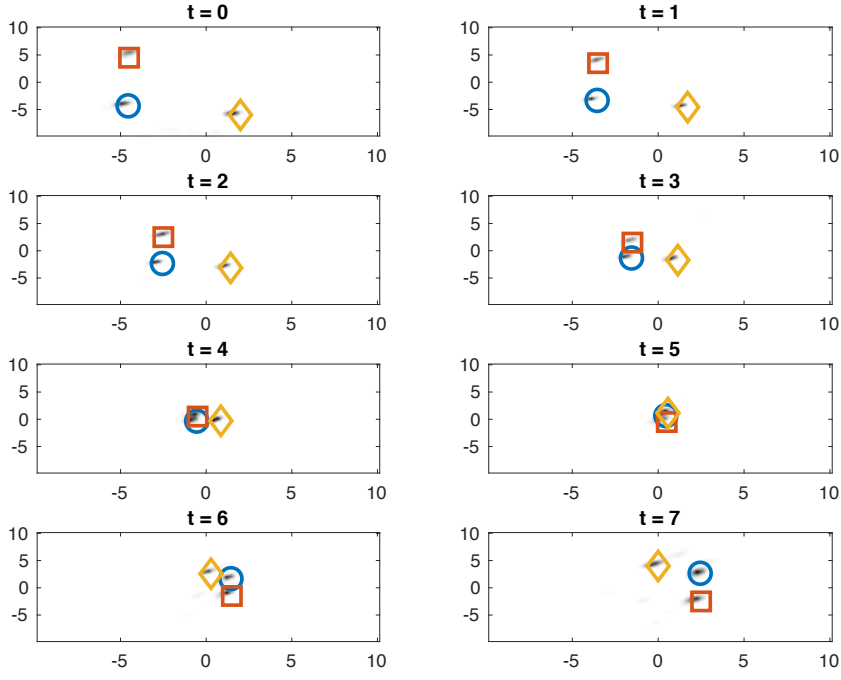

Figure 11: Estimated spatial spectra corresponding to the time points $t = 0, 1, \ldots, \mathcal{T}$, with $\mathcal{T} = 7$, for the barycenter tracking problem. Note that the ground truth locations have been superimposed, using the same legend as in Figure 10.

measure signals generated by three moving sources. The scenario is illustrated in Figure 10. Here, the start location of each source is indicated by an asterisk. As can be seen, the trajectories of the targets all intersect at some point. Note that one of the ULAs has been rotated slightly, as indicated by the orientational arrows, motivating the use of a barycenter formulation in order to form an estimate of the spatial spectrum. At $\mathcal{T} + 1$ time points t, $t = 0, 1, \ldots, \mathcal{T}$, with $\mathcal{T} = 7$, we collect 100 signal snapshots which are used to estimate the covariance matrices of the two arrays using the sample covariance matrix. The target signals are modeled as independent Gaussian sources, and spatially white Gaussian noise is added to the sensors, with an SNR of 20 dB. In order to model the tracking part, we use a two-dimensional extension of the state space description utilized in

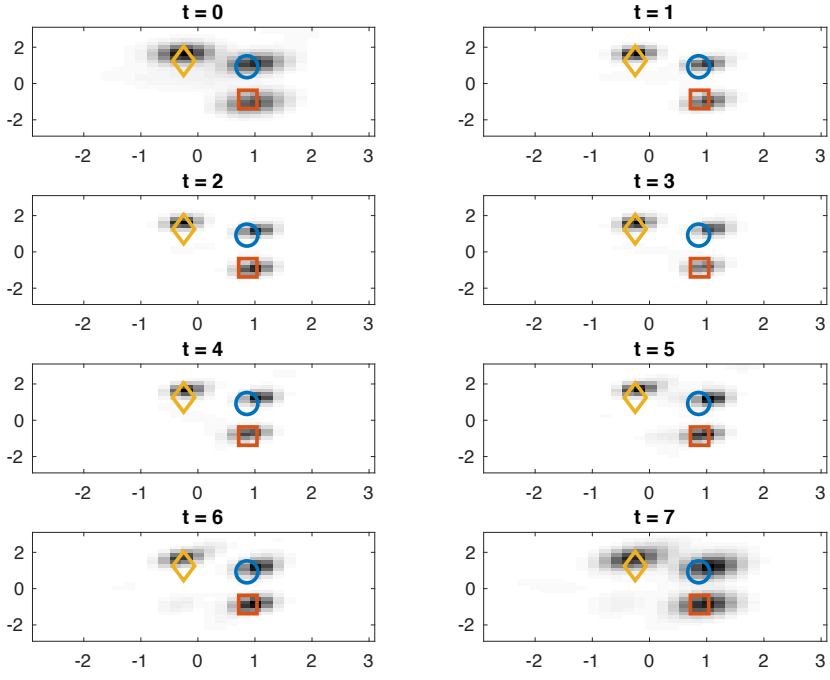

Figure 12: Estimated velocity spectra corresponding to the time points $t = 0, 1, \ldots, \mathcal{T}$, with $\mathcal{T} = 7$, for the barycenter tracking problem. Note that the ground truth locations have been superimposed, using the same legend as in Figure 10.

Section 6.1, where each of the two spatial components is endowed with a velocity state. With this description, the barycenter tracking problem considers transport on a four-dimensional space, i.e., over location and velocity for each spatial dimension. As the cost of transport between the barycenter at the observation times and the corresponding observation marginals, we use the Euclidean cost, i.e., $c(\mathrm{x}_0, \mathrm{x}_1) = \|\mathrm{x}_0 - \mathrm{x}_1\|_2^2$ for $\mathrm{x}_0, \mathrm{x}_1 \in \mathbb{R}^2$. In forming the spectral estimates, we use the discretized version of the problem in (26) using the parameters $\varepsilon = 0.2$, $\gamma = 0.5$, common for all marginals, and $\alpha = 1$. The spatial domain is gridded uniformly with $n_x = 75$ points in each dimension, and the velocity domain is gridded uniformly with $n_v = 30$ points in each dimension.

The resulting estimated spatial spectra are shown in Figure 11. As can be

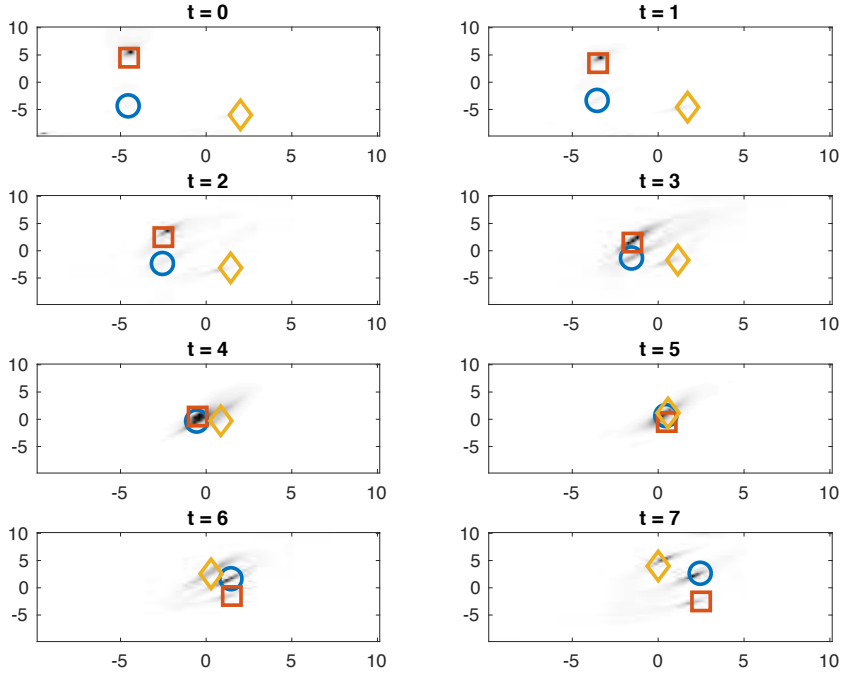

Figure 13: Estimated spatial spectra corresponding to the time points $t = 0, 1, \ldots, \mathcal{T}$, with $\mathcal{T} = 7$, obtained using the non-coherent MVDR estimator. Note that the ground truth locations have been superimposed, using the same legend as in Figure 10.

seen, the modes of the estimate correspond well to the ground truth, taking into account that the array rotation prevents perfect estimates. Note here that the rotation only causes a slight shift in location of the sources, but no spurious estimates. The corresponding estimated velocity spectra are shown in Figure 12. As can be seen, the distribution over velocity details three distinct modes that remain fairly constant over time, corresponding well to the ground truth constant velocity. As comparison, Figure 13 displays the results obtained using the non-coherent MVDR estimator from [56]. Note that this estimator does not take any time dependence into account, and instead forms estimates at each separate time point t using the available pair of array covariance matrices. It may be noted that these estimates display less concentrated estimates, as well as big differences in

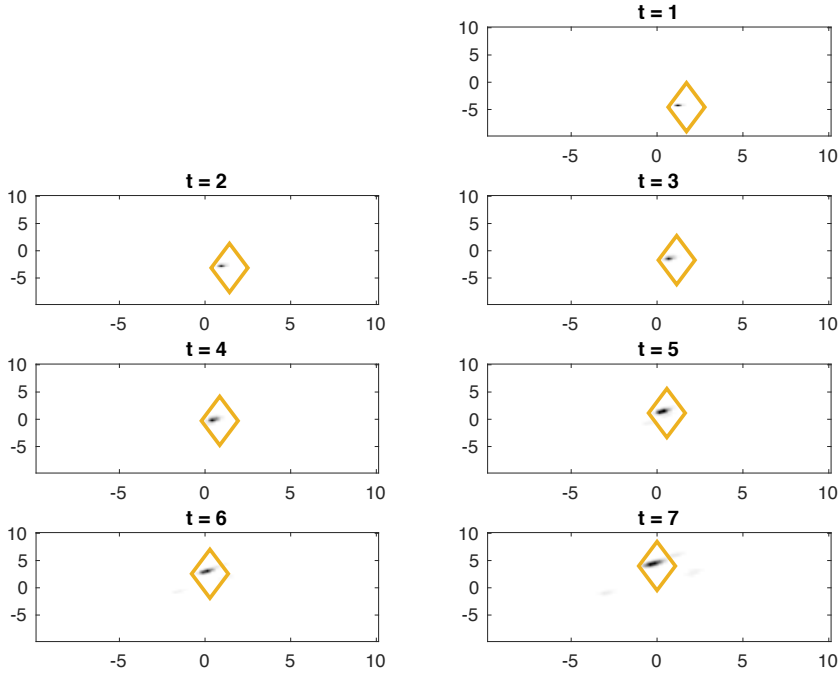

Figure 14: Tracking of the mass of the peak location corresponding to the third target at time $t = 0$.

spectral power.

Finally, in order to illustrate how mass is transported on the 2-D spatial domain, Figure 14 tracks the location of the spectral peak corresponding to the third target along the trajectory, identified by computing the bi-marginal projections corresponding to consecutive barycenters, as detailed in Proposition 2 and Remark 5. Note here that the distribution of mass remains fairly concentrated throughout the trajectory.

6.4.1 Dimensionality and complexity analysis

In this section, we provide a brief dimensionality analysis of the OMT problem used to model the barycenter tracking above, as well an analysis of the computational complexity of solving the dual problem in (36) using Algorithm 2. Recall that the state space consists of a location and velocity component for each

of the two spatial dimensions. Thus, the size of the discrete state space X is $N = (n_x n_v)^d$, where $n_x = 75$ and $n_v = 30$ are the number of grid points on the spatial and velocity domain, respectively, and $d = 2$ is the number of spatial dimensions. Accordingly, the size of the discrete measurement space Y is $\tilde{N} = n_x^d$. As covariance matrices are estimated for $J = 2$ sensor arrays at $\mathcal{T} + 1$ time instances, with $\mathcal{T} = 7$, this implies that the number of elements in the mass transport tensor $\mathbf{M}$ in (34) is

$$N^{(\mathcal{T}+1)}\tilde{N}^{(T+1)J} \approx 4.3 \cdot 10^{113}.$$

This number is larger than the number of particles in the observable universe; solving the primal OMT problem directly, or even constructing a tensor of this size, is thus infeasible. However, the dual problem in (36) is formulated in terms of the dual vectors $\lambda_{(t,j)} \in \mathbb{R}^m$, for $t = 0, \ldots, \mathcal{T}$, and $j = 1, \ldots, J$. Here, $m = 225$ denotes the size of covariance vectors in (33), constructed by stacking the real and imaginary components of the corresponding covariance matrices, excluding the redundant parts resulting from the Hermitian structure. Thus, the number of real variables in the dual problem is given by $(\mathcal{T} + 1)Jm = 3600$, constituting a dramatic complexity reduction as compared to the primal problem.

To analyze the computational complexity for solving the dual problem, we consider the complexity of performing one iteration sweep of Algorithm 2, i.e., one sweep through all index pairs (t, j), for $t = 0, \ldots, \mathcal{T}$, and $j = 1, \ldots, J$, as to update the corresponding dual vectors $\lambda_{(t,j)}$, where we run through j in the inner cycle and t in the outer cycle. It may be noted that for each index pair (t, j), this requires computing the vector $v_{(t,j)}$ in (37), followed by solving (38), as detailed in Theorem 1. The computation of $v_{(t,j)}$ is detailed in the projection in (46) in Proposition 4, disregarding the vector $u_{(t,j)}$. As noted in Section 5.3, computing the set of vectors $v_{(t,j)}$, for a fixed t and $j = 1, 2, \ldots, J$, requires only one multiplication with $K \in \mathbb{R}^{N \times N}$, and $2J$ multiplications with $\tilde{K} \in \mathbb{R}^{N \times n_x^2}$. By exploiting the structures of K and $\tilde{K}$, as described in Remark 4, i.e., the decoupling in the spatial dimensions, the total complexity for performing the update of the left and right factors and computing the set $\left\{v_{(t,j)}\right\}_{j=1}^{J}$, is $\mathcal{O}(d(n_x n_v)^3 + Jdn_x^3 n_v)$. Finding the roots of (38) by Newton's method requires solving a system of linear equations of size $m \times m$ in each inner Newton iteration. However, as indicated in Remark 3, after a few outer Sinkhorn iterations, Newton's method in general converges directly, i.e., it suffices to solve a single system of equations in each outer iteration. Thus, after these initial outer iterations, the complexity for

updating all dual variables $\lambda_{(t,j)}$, for $t = 0, 1, \ldots, \mathcal{T}$, and $j = 1, 2, \ldots, J$, is $\mathcal{O}\left(\mathcal{T}(d(n_x n_v)^3 + Jdn_x^3 n_v + Jm^3)\right)$.

7 Conclusions

In this work, we have proposed a framework for formulating multi-marginal OMT problems for scenarios in which the underlying mass distributions are only indirectly observable, referred to as partial information of the marginals. Motivated by examples from spatial spectral estimation in array processing, we have shown that the proposed formulations may be used for modeling information fusion, as well as for tracking the evolution of mass distributions over time. By leveraging the geometrical properties of OMT, the proposed formulations have been shown to yield robust spectral estimates, as well as allowing for exploiting prior knowledge of underlying dynamics. Also, we have presented computational tools, leading to computationally efficient solution algorithms for the transport problem. Even though the original primal OMT formulation may be prohibitively large, we have shown that by considering dual formulations, as well as exploiting inherent structures in the problem, one may arrive at tractable solvers even in high-dimensional settings.

8 Appendix

In this section, we provide proofs of Propositions 1-4.

Proof of Proposition 1. As to simplify the exposition, let λ and Δ denote the sets $\lambda_0, \ldots, \lambda_{\mathcal{T}}$ and $\Delta_0, \ldots, \Delta_{\mathcal{T}}$, respectively. The Lagrangian of (34), with dual variable λ, is detailed as

$$\begin{aligned} L(\mathbf{M}, \Delta, \lambda) = &\sum_{i_0,\ldots,i_{\mathcal{T}}} \mathbf{C}_{i_0,\ldots,i_{\mathcal{T}}} \mathbf{M}_{i_0,\ldots,i_{\mathcal{T}}} + \varepsilon \mathcal{D}(\mathbf{M}) \\ &+ \sum_{t=0}^{\mathcal{T}} \gamma_t \|\Delta_t\|_2^2 + \sum_{t=0}^{\mathcal{T}} \lambda_t^T (r_t + \Delta_t - G_t P_t(\mathbf{M})). \end{aligned} \tag{47}$$

For fixed λ, the Lagrangian is minimized when the gradients with respect to $\mathbf{M}$

and Δ vanish. For $\mathbf{M}$, this requires

$$\mathbf{C}_{i_0,\ldots,i_\mathcal{T}} + \varepsilon \log \mathbf{M}_{i_0,\ldots,i_\mathcal{T}} - \sum_{t=0}^{\mathcal{T}} \left(\lambda_t^T G_t\right)_{i_t} = 0$$

$$\Leftrightarrow \mathbf{M}_{i_0,\ldots,i_\mathcal{T}} = \exp(-\mathbf{C}_{i_0,\ldots,i_\mathcal{T}}/\varepsilon) \prod_{t=0}^{\mathcal{T}} \exp(G_t^T \lambda_t/\varepsilon).$$

Thus, one may express the mass transport tensor as $\mathbf{M} = \mathbf{K} \odot \mathbf{U}$ for two sensors $\mathbf{K}, \mathbf{U} \in \mathbb{R}^{n^{\mathcal{T}+1}}$, defined as $\mathbf{K} = \exp(-\mathbf{C}/\varepsilon)$ and

$$\mathbf{U} = u_0 \otimes u_1 \otimes \cdots \otimes u_\mathcal{T}, \qquad \text{with } u_t = \exp(G_t^T \lambda_t/\varepsilon).$$

This proves the first statement of the proposition. Further, the Lagrangian (47) is minimized with respect to Δ when

$$2\gamma_t \Delta_t + \lambda_t = 0 \quad \Rightarrow \; \Delta_t = -\frac{1}{2\gamma_t}\lambda_t \quad \text{for } t = 0, 1, \ldots, \mathcal{T}.$$

Thus, the Langrangian (47), when minimized with respect to $\mathbf{M}$ and Δ, becomes

$$\min_{\mathbf{M},\Delta} L(\mathbf{M}, \Delta, \lambda) = \; -\varepsilon \sum_{i_0,\ldots,i_\mathcal{T}} \mathbf{M}_{i_0,\ldots,i_\mathcal{T}} - \sum_{t=0}^{\mathcal{T}} \left(\frac{1}{4\gamma_t} \|\lambda_t\|_2^2 - \lambda_t^T r_t \right),$$

yielding the dual problem

$$\underset{\lambda_0,\ldots,\lambda_\mathcal{T}}{\text{maximize}} \; -\varepsilon \sum_{i_0,\ldots,i_\mathcal{T}} \mathbf{K}_{i_0,\ldots,i_\mathcal{T}} \mathbf{U}_{i_0,\ldots,i_\mathcal{T}} - \sum_{t=0}^{\mathcal{T}} \frac{1}{4\gamma_t} \|\lambda_t\|_2^2 + \sum_{t=0}^{\mathcal{T}} \lambda_t^T r_t.$$

□

The proofs of Proposition 2 - 4, for exploiting tensor structures in order to compute the projections on the marginals, are based on the following two lemmas.

Lemma 1. *Let* $\mathbf{U} = u_0 \otimes u_1 \otimes \cdots \otimes u_\mathcal{T}$ *and suppose that*

$$\langle \mathbf{K}, \mathbf{U} \rangle = w_1^T \operatorname{diag}(u_t) w_2,$$

where w_1 *and* w_2 *are vectors that may depend on* u_ℓ*, for* $\ell \neq t$*. Then,*

$$P_t(\mathbf{K} \odot \mathbf{U}) = w_1 \odot u_t \odot w_2.$$

Proof. Let e_{i_t} denote the unit vector whose element i_t is equal to one, and with all other elements equal to zero. Then, by replacing u_t with $u_t \odot e_{i_t}$ in the expression for $\langle \mathbf{K}, \mathbf{U} \rangle$, we get the element i_t of the projection on the t:th marginal, i.e.,

$$\begin{aligned}
&\langle \mathbf{K}, u_0 \otimes \cdots \otimes u_{t-1} \otimes (u_t \odot e_{i_t}) \otimes u_{t+1} \otimes \cdots \otimes u_{\mathcal{T}} \rangle \\
&= \sum_{\ell_0, \ell_1 \ldots, \ell_{\mathcal{T}}} \mathbf{K}_{\ell_0, \ell_1 \ldots, \ell_{\mathcal{T}}} \left(\prod_{s=0}^{t-1} (u_s)_{\ell_s} \right) (u_t \odot e_{i_t})_{\ell_t} \left(\prod_{s=t+1}^{\mathcal{T}} (u_s)_{\ell_s} \right) \\
&= \sum_{\substack{\ell_0, \ldots, \ell_{t-1} \\ \ell_{t+1}, \ldots, \ell_{\mathcal{T}}}} \mathbf{K}_{\ell_0, \ell_1 \ldots, \ell_{\mathcal{T}}} \left(\prod_{s=0}^{t-1} (u_s)_{\ell_s} \right) (u_t)_{i_t} \left(\prod_{s=t+1}^{\mathcal{T}} (u_s)_{\ell_s} \right) \\
&= (P_t(\mathbf{K} \odot \mathbf{U}))_{i_t}.
\end{aligned}$$

It may be noted that in the second line, a term in the sum is only non-zero if $i_t = \ell_t$, yielding the second equality. By the assumption in the lemma, we get

$$(P_t(\mathbf{K} \odot \mathbf{U}))_{i_t} = w_1^T \operatorname{diag}(u_t \odot e_{i_t}) w_2 = (w_1)_{i_t} (u_t)_{i_t} (w_2)_{i_t},$$

from which the result follows. □

Lemma 2. *Let* $\mathbf{U} = u_0 \otimes u_1 \otimes \cdots \otimes u_{\mathcal{T}}$ *and suppose that*

$$\langle \mathbf{K}, \mathbf{U} \rangle = w_1^T \operatorname{diag}(u_{t_1}) W_2 \operatorname{diag}(u_{t_2}) w_3,$$

where $w_1, w_3 \in \mathbb{R}^n$ *and* $W_2 \in \mathbb{R}^{n \times n}$ *may depend on* u_ℓ*, for* $\ell \notin \{t_1, t_2\}$*. Then,*

$$P_{t_1, t_2}(\mathbf{K} \odot \mathbf{U}) = \operatorname{diag}(w_1 \odot u_{t_1}) W_2 \operatorname{diag}(u_{t_2} \odot w_3).$$

Proof. Analogously to the proof in Lemma 1, we may express the bi-marginal projections as

$$\begin{aligned}
(P_{t_1, t_2}(\mathbf{K} \odot \mathbf{U}))_{i_{t_1}, i_{t_2}} = \langle \mathbf{K}, u_0 \otimes \cdots \otimes u_{t_1 - 1} \otimes (u_{t_1} \odot e_{i_{t_1}}) \otimes u_{t_1 + 1} \otimes \cdots \\
\cdots \otimes u_{t_2 - 1} \otimes (u_{t_2} \odot e_{i_{t_2}}) \otimes u_{t_2 + 1} \otimes \cdots \otimes u_{\mathcal{T}} \rangle.
\end{aligned}$$

Hence, by the assumption,

$$\begin{aligned}
(P_{t_1, t_2}(\mathbf{K} \odot \mathbf{U}))_{i_{t_1}, i_{t_2}} &= w_1^T \operatorname{diag}(u_{t_1} \odot e_{i_{t_1}}) W_2 \operatorname{diag}(u_{t_2} \odot e_{i_{t_2}}) w_3 \\
&= (w_1)_{i_{t_1}} (u_{t_1})_{i_{t_1}} (W_2)_{i_{t_1}, i_{t_2}} (u_{t_2})_{i_{t_2}} (w_3)_{i_{t_2}},
\end{aligned}$$

and thus the result follows.

□

Proof of Proposition 2. First, note that due to the assumption of a sequential cost, we have

$$\mathbf{K}_{i_0,\ldots,i_T} = \prod_{t=1}^{\mathcal{T}} K_{i_{t-1},i_t},$$

for the tensor $\mathbf{K} = \exp(-\mathbf{C}/\varepsilon)$ and matrix $K = \exp(-C/\varepsilon)$. Therefore,

$$\begin{aligned}
\langle \mathbf{K}, \mathbf{U} \rangle &= \sum_{i_0,i_1,\ldots,i_{\mathcal{T}}} \left(\prod_{t=1}^{\mathcal{T}} K_{i_{t-1},i_t} \right) \prod_{t=0}^{\mathcal{T}} (u_t)_{i_t} \\
&= \sum_{i_0,i_1,\ldots,i_{\mathcal{T}}} (u_0)_{i_0} \prod_{s=1}^{\mathcal{T}} (K \operatorname{diag}(u_t))_{i_{t-1},i_t} \\
&= u_0^T K \operatorname{diag}(u_1) K \ldots K \operatorname{diag}(u_{\mathcal{T}-1}) K u_{\mathcal{T}} \\
&= \mathbf{1}^T \operatorname{diag}(u_0) K \operatorname{diag}(u_1) K \ldots K \operatorname{diag}(u_{\mathcal{T}-1}) K \operatorname{diag}(u_{\mathcal{T}}) \mathbf{1}.
\end{aligned}$$

Thus, $\langle \mathbf{K}, \mathbf{U} \rangle$ may be written as in Lemma 1 with

$$\begin{aligned}
w_1 &= \left(\mathbf{1}^T \operatorname{diag}(u_0) K \operatorname{diag}(u_1) K \ldots K \operatorname{diag}(u_{t-1}) K \right)^T, \\
w_2 &= K \operatorname{diag}(u_{t+1}) K \ldots K \operatorname{diag}(u_{\mathcal{T}-1}) K \operatorname{diag}(u_{\mathcal{T}}) \mathbf{1},
\end{aligned}$$

and hence,

$$\begin{aligned}
P_t(\mathbf{K} \odot \mathbf{U}) = &\left(u_0^T K \operatorname{diag}(u_1) K \ldots K \operatorname{diag}(u_{t-1}) K \right)^T \odot u_t \\
&\odot \left(K \operatorname{diag}(u_{t+1}) K \ldots K \operatorname{diag}(u_{\mathcal{T}-1}) K u_{\mathcal{T}} \right).
\end{aligned}$$

Moreover, in order to derive an expression for $P_{t_1,t_2}(\mathbf{K} \odot \mathbf{U})$, note that $\langle \mathbf{K}, \mathbf{U} \rangle$ can be written as in Lemma 2 with

$$\begin{aligned}
w_1 &= \left(\mathbf{1}^T \operatorname{diag}(u_0) K \operatorname{diag}(u_1) K \ldots K \operatorname{diag}(u_{t_1-1}) K \right)^T \\
W_2 &= K \operatorname{diag}(u_{t_1+1}) K \ldots K \operatorname{diag}(u_{t_2-1}) K \\
w_3 &= K \operatorname{diag}(u_{t_2+1}) K \ldots K \operatorname{diag}(u_{\mathcal{T}-1}) K \operatorname{diag}(u_{\mathcal{T}}) \mathbf{1},
\end{aligned}$$

and hence the expression (40) follows.

□

Proof of Proposition 3. Due to the structure of the cost tensor $\mathbf{C}$, we have

$$\mathbf{K}_{i_0,\dots,i_J} = \prod_{j=1}^{J} K_{i_0,i_j},$$

for the tensor $\mathbf{K} = \exp(-\mathbf{C}/\varepsilon)$ and the matrix $K = \exp(-C/\varepsilon)$. Thus, one may write

$$\begin{aligned}
\langle \mathbf{K}, \mathbf{U} \rangle &= \sum_{i_0,i_1,\dots,i_J} \left(\prod_{\ell=1}^{J} K_{i_0,i_\ell} \right) \prod_{\ell=0}^{J} (u_\ell)_{i_\ell} \\
&= \sum_{i_0} (u_0)_{i_0} \sum_{i_1,\dots,i_J} \prod_{\ell=1}^{J} K_{i_0,i_\ell} (u_\ell)_{i_\ell} \\
&= \sum_{i_0} (u_0)_{i_0} \prod_{\ell=1}^{J} (K u_\ell)_{i_0} \\
&= u_0^T \left(\bigodot_{\ell=1}^{J} K u_\ell \right).
\end{aligned}$$

It may be noted that this may be expressed as

$$\langle \mathbf{K}, \mathbf{U} \rangle = \mathbf{1}^T \operatorname{diag}(u_0) \left(\bigodot_{\ell=1}^{J} K u_\ell \right),$$

as well as

$$\begin{aligned}
\langle \mathbf{K}, \mathbf{U} \rangle &= (K u_j)^T \left(u_0 \odot \bigodot_{\substack{\ell=1 \\ \ell \neq j}}^{J} K u_\ell \right) \\
&= \mathbf{1}^T \operatorname{diag}(u_j) K^T \left(u_0 \odot \bigodot_{\substack{\ell=1 \\ \ell \neq j}}^{J} K u_\ell \right).
\end{aligned}$$

Applying Lemma 1 yields the expressions (41) and (42) for $P_j(\mathbf{K} \odot \mathbf{U})$ for $j = 0, \dots, J$. Alternatively, one may rewrite $\langle \mathbf{K}, \mathbf{U} \rangle$ as

$$\langle \mathbf{K}, \mathbf{U} \rangle = \mathbf{1}^T \operatorname{diag}(u_0) \operatorname{diag} \left(\bigodot_{\substack{\ell=1 \\ \ell \neq j}}^{J} K u_\ell \right) K \operatorname{diag}(u_j) \mathbf{1},$$

and

$$\langle \mathbf{K}, \mathbf{U} \rangle = \mathbf{1}^T \operatorname{diag}(u_{j_1}) K^T \operatorname{diag}\left(u_0 \odot \bigodot_{\substack{\ell=1 \\ \ell \neq j_1, j_2}}^{J} K u_\ell \right) K \operatorname{diag}(u_{j_2}) \mathbf{1}.$$

With Lemma 2, this leads to the expressions (43) and (44) for the couplings $P_{j_1,j_2}(\mathbf{K} \odot \mathbf{U})$ for $j_1, j_2 = 0, \ldots, J$.

□

Proof of Proposition 4. Recall the definition of the set $\Lambda = \{(t,j) \,|\, t \in \{0,1,\ldots,\mathcal{T}\}, j \in \{0,1\ \ldots,J\}\}$. The structure of the cost tensor $\mathbf{C}$ then implies that each element of the tensor $\mathbf{K} = \exp(-\mathbf{C}/\varepsilon)$ may be expressed as

$$\mathbf{K}_{(i_{(t,j)} | (t,j) \in \Lambda)} = \left(\prod_{t=1}^{\mathcal{T}} K_{i_{(t-1,0)} i_{(t,0)}} \right) \prod_{t=0}^{\mathcal{T}} \prod_{j=1}^{J} \tilde{K}_{i_{(t,0)} i_{(t,j)}},$$

with the matrices defined as $K = \exp(-C/\varepsilon)$ and $\tilde{K} = \exp(-\tilde{C}/\varepsilon)$. Furthermore, the elements of the tensor $\mathbf{U}$ are given by

$$\mathbf{U}_{(i_{(t,j)} | (t,j) \in \Lambda)} = \prod_{t=0}^{\mathcal{T}} \prod_{j=0}^{J} (u_{(t,j)})_{i_{(t,j)}}.$$

Therefore,

$$\begin{aligned}
\langle \mathbf{K}, \mathbf{U} \rangle &= \sum_{\substack{i_{(s,\ell)} \\ (s,\ell) \in \Lambda}} \left(\prod_{t=1}^{\mathcal{T}} K_{i_{(t-1,0)} i_{(t,0)}} \right) \left(\prod_{t=0}^{\mathcal{T}} \prod_{j=1}^{J} \tilde{K}_{i_{(t,0)} i_{(t,j)}} \right) \prod_{t=0}^{\mathcal{T}} \prod_{j=0}^{J} (u_{(t,j)})_{i_{(t,j)}} \\
&= \sum_{\substack{i_{(s,0)} \\ s=0,1,\ldots,\mathcal{T}}} \left(\prod_{t=1}^{\mathcal{T}} K_{i_{(t-1,0)} i_{(t,0)}} \right) \prod_{t=0}^{\mathcal{T}} \left(\sum_{\substack{i_{(t,\ell)} \\ \ell=1,2,\ldots,J}} \left(\prod_{j=1}^{J} \tilde{K}_{i_{(t,0)} i_{(t,j)}} \right) \prod_{j=0}^{J} (u_{(t,j)})_{i_{(t,j)}} \right).
\end{aligned}$$

Note that the last sum is a projection of the type (41) in Proposition 3, and we denote it as

$$\sum_{\substack{i_{(t,\ell)} \\ \ell=1,2,\ldots,J}} \left(\prod_{j=1}^{J} \tilde{K}_{i_{(t,0)} i_{(t,j)}} \right) \prod_{j=0}^{J} (u_{(t,j)})_{i_{(t,j)}} = \left(u_{(t,0)} \odot \bigodot_{j=1}^{J} \left(\tilde{K} u_{(t,j)} \right) \right)_{i_{(t,0)}}$$
$$= (p_t)_{i_{(t,0)}}.$$

Hence, using the proof of Proposition 2, one may write

$$\begin{aligned} \langle \mathbf{K}, \mathbf{U} \rangle =& p_0^T K \operatorname{diag}(p_1) K \ldots K \operatorname{diag}(p_{\mathcal{T}-1}) K p_{\mathcal{T}} \\ =& p_0^T K \ldots K \operatorname{diag}(p_{t-1}) K \operatorname{diag}(u_{(t,0)}) \operatorname{diag} \left(\bigodot_{j=1}^{J} \left(\tilde{K} u_{(t,j)} \right) \right) \\ & \cdot K \operatorname{diag}(p_{t+1}) K \ldots K p_{\mathcal{T}}. \end{aligned}$$

Thus, the projections on the central marginals, corresponding to index $(t,0)$ for $t = 0, 1, \ldots, \mathcal{T}$, may be computed as in Proposition 2, yielding

$$\begin{aligned} P_{(t,0)}(\mathbf{K} \odot \mathbf{U}) = & \left(p_0^T K \operatorname{diag}(p_1) K \ldots \operatorname{diag}(p_{t-1}) K \right)^T \odot p_t \\ & \odot \left(K \operatorname{diag}(p_{t+1}) \ldots K \operatorname{diag}(p_{\mathcal{T}-1}) K p_{\mathcal{T}} \right). \end{aligned}$$

For the other marginals, i.e., for (t,j) such that $j > 0$, one may express $\langle \mathbf{K}, \mathbf{U} \rangle$ similarly as in the proof of (42) of Proposition 3, i.e.,

$$\begin{aligned} \langle \mathbf{K}, \mathbf{U} \rangle =& \mathbf{1}^T \operatorname{diag}(u_{(t,j)}) \tilde{K}^T \Big((p_0^T K \operatorname{diag}(p_1) K \ldots \operatorname{diag}(p_{t-1}) K)^T \\ & \odot (p_t ./ (K u_{(t,j)})) \odot (K \operatorname{diag}(p_{t+1}) \ldots K \operatorname{diag}(p_{\mathcal{T}-1}) K p_{\mathcal{T}}) \Big), \end{aligned}$$

which, with Lemma 1, yields the expression (46).

□

References

[1] P. Stoica and R. Moses, *Spectral Analysis of Signals*, Prentice Hall, Upper Saddle River, N.J., 2005.

[2] Y. Hu and P. C. Loizou, "A Perceptually Motivated Approach for Speech Enhancement," *IEEE Speech Audio Process.*, vol. 11, no. 5, pp. 457–465, Sept. 2003.

[3] A. N. Amini, E. Ebbini, and T. T. Georgiou, "Noninvasive estimation of tissue temperature via high-resolution spectral analysis techniques," *IEEE Trans. Bio. Eng.*, vol. 52, no. 2, pp. 221–228, Feb. 2005.

[4] S. Mierisová and M. Ala-Korpela, "MR spectroscopy quantitation: a review of frequency domain methods," *NMR in Biomedicine*, vol. 14, no. 4, pp. 247–259, 2001.

[5] H. L. Van Trees, *Detection, Estimation, and Modulation Theory: Radar-Sonar Signal Processing and Gaussian Signals in Noise*, Krieger Publishing Co., Inc., 1992.

[6] S. Kullback and R. A. Leibler, "On Information and Sufficiency," *The Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79–86, March 1951.

[7] R. M. Gray, A. Buzo, A.. H. Gray Jr, and Y. Matsuyama, "Distortion measures for speech processing," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 28, no. 4, pp. 367–376, 1980.

[8] A. Rényi, "On measures of entropy and information," in *Proc. 4th Berkeley Sym. Mathematics of Statistics and Probability*, 1961, pp. 547–561.

[9] P. Stoica, P. Babu, and J. Li, "SPICE : a novel covariance-based sparse estimation method for array processing," *IEEE Trans. Signal Process.*, vol. 59, no. 2, pp. 629 –638, Feb. 2011.

[10] J. Swärd, S. I. Adalbjörnsson, and A. Jakobsson, "Generalized Sparse Covariance-based Estimation," *Elsevier Signal Processing*, vol. 143, pp. 311–319, February 2018.

[11] T. T. Georgiou, J. Karlsson, and M. S. Takyar, "Metrics for power spectra: an axiomatic approach," *IEEE Trans. Signal Process.*, vol. 57, no. 3, pp. 859–867, Mar. 2009.

[12] C. Villani, *Optimal transport: old and new*, Springer Science & Business Media, 2008.

[13] M. Beiglböck, P. Henry-Labordère, and F. Penker, "Model-independent bounds for option prices - a mass transport approach," *Finance and Stochastics*, vol. 17, no. 3, pp. 477–501, 2013.

[14] A. Dominitz and A. Tannenbaum, "Texture Mapping via Optimal Mass Transport," *IEEE Trans. Vis. Comput. Graphics*, vol. 16, no. 3, pp. 419–433, 2010.

[15] M. A. Schmitz, M. Heitz, N. Bonneel, F. Ngolè, D. Coeurjolly, M. Cuturi, G. Peyré, and J-L. Starck, "Wasserstein Dictionary Learning: Optimal Transport-Based Unsupervised Nonlinear Dictionary Learning," *SIAM J. Imaging Sci.*, vol. 11, no. 1, pp. 643–678, 2018.

[16] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein GAN," *arXiv preprint arXiv:1701.07875*, 2017.

[17] S. Kolouri, S. R. Park, M. Thorpe, D. Slepcev, and G. K. Rohde, "Optimal Mass Transport: Signal processing and machine-learning applications," *IEEE Signal Process. Mag.*, vol. 34, no. 4, pp. 43–59, July 2017.

[18] M. Cuturi, "Sinkhorn distances: Lightspeed computation of optimal transport," *Proc. Adv. Neural Inf. Process. Syst.*, pp. 2292–2300, 2013.

[19] G. Peyré and M. Cuturi, "Computational optimal transport," *Foundations and Trends® in Machine Learning*, vol. 11, no. 5-6, pp. 355–607, 2019.

[20] F. Elvander, I. Haasler, A. Jakobsson, and J. Karlsson, "Tracking and Sensor Fusion in Direction of Arrival Estimation Using Optimal Mass Transport," in *26th European Signal Processing Conference*, Rome, Italy, Sep. 3-7 2018, pp. 1617–1621.

[21] F. Elvander, I. Haasler, A. Jakobsson, and J. Karlsson, "Non-Coherent Sensor Fusion via Entropy Regularized Optimal Mass Transport," in *Proc. 44th IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, Brighton, UK, May 13-17 2019, pp. 4415–4419.

[22] R. J. McCann, "A convexity principle for interacting gases," *Advances in mathematics*, vol. 128, no. 1, pp. 153–179, 1997.

[23] S. Angenent, S. Haker, and A. Tannenbaum, "Minimizing Flows for the Monge-Kantorovich Problem," *SIAM J. Math. Anal.*, vol. 35, no. 1, pp. 61–97, 2003.

[24] E. Simou and P. Frossard, "Graph Signal Representation with Wasserstein Barycenters," in *Proc. 44th IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, Brighton, UK, May 13-17 2019, pp. 5386–5390.

[25] X. Jiang, Z. Q. Luo, and T. T. Georgiou, "Geometric Methods for Spectral Analysis," *IEEE Trans. Signal Process.*, vol. 60, no. 3, pp. 1064–1074, Mar. 2012.

[26] L. Ning, T. T. Georgiou, and A. Tannenbaum, "On Matrix-Valued Monge-Kantorovich Optimal Mass Transport," *IEEE Trans. Autom. Control*, vol. 60, no. 2, pp. 373–382, 2015.

[27] L. Ning, X. Jiang, and T. T. Georgiou, "On the Geometry of Covariance Matrices," *IEEE Signal Process. L*, vol. 20, no. 8, pp. 787–790, Aug. 2013.

[28] F. Elvander, A. Jakobsson, and J. Karlsson, "Interpolation and Extrapolation of Toeplitz Matrices via Optimal Mass Transport," *IEEE Trans. Signal. Process*, vol. 66, no. 20, pp. 5285 – 5298, Oct. 2018.

[29] K. Yamamoto, Y. Chen, L. Ning, T. T. Georgiou, and A. Tannenbaum, "Regularization and Interpolation of Positive Matrices," *IEEE Trans. Autom. Control*, vol. 63, no. 4, pp. 1208–1212, 2018.

[30] J. D. Benamou, G. Carlier, M. Cuturi, L. Lenna, and G. Peyré, "Iterative Bregman Projections for Regularized Transportation Problems," *SIAM Journal of Scientific Computing*, vol. 37, no. 2, pp. 1111–1138, 2015.

[31] J. Karlsson and A. Ringh, "Generalized Sinkhorn iterations for regularizing inverse problems using optimal mass transport," *SIAM Journal on Imaging Sciences*, vol. 10, no. 4, pp. 1935–1962, 2017.

[32] Y. Chen and J. Karlsson, "State Tracking of Linear Ensembles via Optimal Mass Transport," *IEEE Control Systems Letters*, vol. 2, no. 2, pp. 260–265, Apr 2018.

[33] J. Lellmann, D. A. Lorenz, C. Schönlieb, and T. Valkonen, "Imaging with Kantorovich–Rubinstein Discrepancy," *SIAM J. Imaging Sci.*, vol. 7, no. 4, pp. 2833–2859, 2014.

[34] J. Adler, A. Ringh, O. Öktem, and J. Karlsson, "Learning to solve inverse problems using Wasserstein loss," *arXiv preprint arXiv:1710.10898*, 2017.

[35] P. del Aguila Pla and J. Jaldén, "Cell Detection by Functional Inverse Diffusion and Non-negative Group Sparsity - Part II: Proximal Optimization and Performance Evaluation," *IEEE Trans. Signal Process.*, vol. 66, no. 20, pp. 5422–5437, 2018.

[36] Y. Chen, T. T. Georgiou, and M. Pavon, "Optimal Transport Over a Linear Dynamical System," *IEEE Trans. Autom. Control*, vol. 62, no. 5, pp. 2137–2152, May 2017.

[37] T. T. Georgiou, "Solution of the general moment problem via a one-parameter imbedding," *IEEE Trans. Autom. Control*, vol. 50, no. 6, pp. 811–826, 2005.

[38] D. H. Johnson and D. E. Dudgeon, *Array Signal Processing: Concepts and Techniques*, Prentice Hall, Englewood Cliffs, N.J., 1993.

[39] B. Pass, "Multi-marginal optimal transport: theory and applications," *ESAIM: Mathematical Modelling and Numerical Analysis*, vol. 49, no. 6, pp. 1771–1790, 2015.

[40] I. Abraham, R. Abraham, M. Bergounioux, and G. Carlier, "Tomographic Reconstruction from a Few Views: A Multi-Marginal Optimal Transport Approach," *Applied Mathematics and Optimization*, vol. 75, no. 1, pp. 55–73, 2017.

[41] Y. Brenier, "Generalized solutions and hydrostatic approximation of the Euler equations," *Physica D: Nonlinear Phenomena*, vol. 237, pp. 1982–1988, 2008.

[42] L. M. Kaplan, "Global node selection for localization in a distributed sensor network," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 42, no. 1, pp. 113–135, Jan. 2006.

[43] D. Li, K. D. Wong, Y. H. Hu, and A. M. Sayeed, "Detection, classification, and tracking of targets," *IEEE Signal Process. Mag.*, vol. 19, no. 2, pp. 17–29, Mar. 2002.

[44] O. Pele and M. Werman, "Fast and robust Earth Mover's Distances," in *IEEE 12th Int. Conf. on Comp. Vis.*, Kyoto, Japan, Sep. 29 - Oct. 2 2009, pp. 460–467.

[45] C. Léonard, "A survey of the Schrödinger problem and some of its connections with optimal transport," *Discrete & Continuous Dynamical Systems - A*, vol. 34, no. 4, pp. 1533–1574, 2014.

[46] Y. Chen, T.T. Georgiou, and M. Pavon, "On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint," *Journal of Optimization Theory and Applications*, vol. 169, no. 2, pp. 671–691, 2016.

[47] R. Sinkhorn, "Diagonal Equivalence to Matrices with Prescribed Row and Column Sums," *The American Mathematical Monthly*, vol. 74, no. 4, pp. 402–405, 1967.

[48] J. Lorenz and L. Franklin, "On the scaling of multidimensional matrices," *Linear Algebra and its Applications*, vol. 114-115, pp. 717–715, 1989.

[49] J. C. Bezdek and R. J. Hathaway, "Two new convergence results for alternating optimization," in *Computational Intelligence: The Experts Speak*, pp. 149–164. John Wiley and Sons Inc., 2003.

[50] L. Chizat, G. Peyré, B. Schmitzer, and F. X. Vialard, "Scaling algorithms for unbalanced optimal transport problems," *Mathematics of Computation*, vol. 87, no. 314, pp. 2563–2609, 2018.

[51] I. Haasler, A. Ringh, Y. Chen, and J. Karlsson, "Estimating ensemble flows on a hidden Markov chain," in *Proceedings of the 58th IEEE Conference on Decision and Control*, 2019.

[52] N. Vervliet, O. Debals, L. Sorber, M. Van Barel, and L. De Lathauwer, "Tensorlab 3.0," Mar. 2016, Available online.

[53] J. Capon, "High Resolution Frequency Wave Number Spectrum Analysis," *Proc. IEEE*, vol. 57, pp. 1408–1418, 1969.

[54] Auditec, "Auditory Tests (Revised), Compact Disc, Auditec, St. Louis," St. Louis, 1997.

[55] E. De Sena, N. Antonello, M. Moonen, and T. van Waterschoot, "On the Modeling of Rectangular Geometries in Room Acoustic Simulations," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 23, no. 4, pp. 774–786, Apr. 2015.

[56] D. W. Rieken and D. R. Fuhrmann, "Generalizing MUSIC and MVDR for Multiple Noncoherent Arrays," *IEEE Trans. Signal Process.*, vol. 52, no. 9, pp. 2396–2406, Sep. 2004.

[57] J-A. Luo, K. Yu, Z. Wang, and Y-H. Hu, "Passive source localization from array covariance matrices via joint sparse representations," *Neurocomputing*, vol. 270, pp. 82–90, 2017.

[58] W. Suleiman, P. Parvazi, M. Pesavento, and A. M. Zoubir, "Non-Coherent Direction-of-Arrival Estimation Using Partly Calibrated Arrays," *IEEE Trans. Signal Process.*, vol. 66, no. 21, pp. 5776–5788, Nov. 2018.

[59] C. M. See and A. B. Gershman, "Direction-of-arrival estimation in partly calibrated subarray-based sensor arrays," *IEEE Trans. Signal Process.*, 2004.

[60] D. Xie, J. Huang, and H. Ge, "Localization of near-field sources with partly calibrated subarray-based array," in *5th IEEE Conf. on Industrial Electronics and Applications*, 2010, pp. 1758–1761.

[61] A. Adler and M. Wax, "Direct Localization by Partly Calibrated Arrays: A Relaxed Maximum Likelihood Solution," in *27th European Signal Processing Conference*, A Coruna, Spain, 2019.

[62] S. A. Elkader, A. B. Gershman, and K. M. Wong, "Rank reduction direction-of-arrival estimators with an improved robustness against subarray orientation errors," *IEEE Trans. Signal Process.*, vol. 54, no. 5, pp. 1951–1955, 2006.

C

Paper C

Defining Fundamental Frequency for Almost Harmonic Signals

Filip Elvander and Andreas Jakobsson

Centre for Mathematical Sciences, Lund University, Lund, Sweden

Abstract

In this work, we consider the modeling of signals that are almost, but not quite, harmonic, i.e., composed of sinusoids whose frequencies are close to being integer multiples of a common frequency. Typically, in applications, such signals are treated as perfectly harmonic and one attempts to estimate their fundamental frequency, despite the signals not actually being periodic. Herein, we provide three different definitions of a concept of fundamental frequency for such inharmonic signals and study the implications of the different choices for modeling and estimation. We show that one of the definitions corresponds to a misspecified modeling scenario, and provides a theoretically solid benchmark for analyzing the behavior of estimators derived under a perfectly harmonic assumption. The second definition stems from optimal mass transport theory and yields a robust and easily interpretable concept of fundamental frequency based on the signals' spectral properties. The third definition interprets the inharmonic signal as an observation of a randomly perturbed harmonic signal. This allows for computing a hybrid information theoretical bound on estimation performance, as well as for finding an estimator attaining the bound. The theoretical findings are illustrated using numerical examples.

Key words: Inharmonicity, fundamental frequency estimation, misspecified models, optimal mass transport

1 Introduction

Signals that may be well modeled as superpositions of harmonically related sinusoids appear in a wide variety of fields. Such signals, often referred to as pitches, are, for instance, used in speech processing for modeling the voiced part of human speech [1], in music information retrieval for extracting musical melodies [2], for monitoring and fault detection of industrial machinery [3], or assessing the state of human diseases [4]. In such applications, the most commonly employed signal feature is the fundamental frequency, or pitch, corresponding to the reciprocal of the signal period. Due to this, considerable effort has been directed towards developing estimators that are statistically, as well as computationally, efficient [5], with more recent contributions extending the estimation problem to signals containing multiple harmonic structures (see, e.g., [6]). However, in some cases, the assumed harmonic structure is only approximate, i.e., there exists no fundamental frequency such that the frequencies of the sinusoidal are all integer multiples of it, or, equivalently, the signal is not periodic. Such signals are referred to as inharmonic and may be found, e.g., in the sound produced by stringed musical instruments [7] and to some extent the human voice [8]. In the former case, there exists parametric models for some instruments based on their physical properties, describing the precise deviation from a perfect harmonic structure [7]. This allows for the formulation of efficient estimators of the signal parameters [9, 10], as well for deriving information theoretical bounds on estimation performance, such as the Cramér-Rao lower bound (CRLB) [11]. However, when no apparent structure of the inharmonicity is known, it is less clear how to efficiently estimate the frequency content of such signals. Intuitively, one should be able to achieve better estimation performance than what is possible under the assumption of an unstructured model, i.e., when there is no relation between the frequencies of the sinusoidal components. Although one may use robust methods [9] or simply resort to applying estimators derived for the perfectly harmonic case, it is not clear what the relevant bounds on estimation performance are. In fact, when applying a harmonic estimator to an inharmonic signal, it is seldom stated what quantity one actually aims to estimate. This is the question this work aims to answer: what does the concept of a pitch mean for inharmonic signals? We provide three possible definitions, emanating from three different views of the inharmonic signal, and study the implications of the different choices.

Firstly, we follow a traditional approach and propose to define the pitch of inharmonic signals via the best ℓ_2 approximation. We show that for the case

of Gaussian additive noise, this corresponds to the best approximation in the Kullback-Leibler sense, allowing for utilizing the framework of misspecified estimation and for interpreting the pitch definition as a pseudo-true fundamental frequency (see, e.g., [12]). Furthermore, this allows us to use the misspecified CRLB (MCRLB) [13] as a bound on estimation performance, which we derive for the inharmonic model. This bound is asymptotically tight and attained by the misspecified maximum likelihood estimator (MLE). Thus, this definition formalizes the inherent assumptions behind applying harmonic estimators to inharmonic signals. However, although practically useful, we show that this definition is problematic for signals with a long time duration, as well as depends on parameters that may be considered nuisance.

As an alternative, we consider approximating the spectral representation of infinite-length versions of the inharmonic signal. Building on the concept of optimal mass transport (OMT) (see, e.g., [14]), we consider the harmonic spectrum closest in the OMT sense, yielding a definition of the fundamental frequency. We show that this definition has some attractive properties, such as stability to small perturbations, as well as an intuitive appeal. Furthermore, when using the definition as a basis for estimating the pitch, the resulting estimator allows for a closed-form expression for the asymptotic variance, and, in the case of perfect harmonicity and white Gaussian noise, has the same asymptotical performance as the MLE. In this case, the resulting estimator corresponds asymptotically to an estimator formed using the extended invariance principle (EXIP) [15], fitting a set of unstructured frequency estimates to a perfectly harmonic structure. The EXIP concept extends the invariance principle for ML estimation to the case when the variable transformation is not bijective and has earlier been successfully applied in, e.g., array processing problems in the presence of calibration errors [16], as well as for pitch estimation [17] under the assumption of perfect harmonicity.

Lastly, we consider modeling inharmonic signals within a stochastic framework, wherein the noiseless signal is regarded as an observation of a random variable. Specifically, we model deviations from the perfectly harmonic model as zero-mean random variables, allowing for interpreting the pitch as an expectation. As the resulting model contains both deterministic and random parameters, we derive the hybrid CRLB (HCRLB) as a lower bound on estimation mean squared error (MSE). We also derive an easily implementable hybrid ML/maximum a posteriori (MAP) estimator that asymptotically attains the resulting HCRLB.

We compare and contrast the three definitions with each other, highlighting

their relative merits and applicability, as well as provide numerical illustrations of the derived estimators and proposed bounds.

2 Almost harmonic signals

Consider the noiseless signal model[1]

$$x_t = \sum_{k=1}^{K} \tilde{r}_k e^{i\tilde{\varphi}_k + i\tilde{\omega}_k t}, \tag{1}$$

for $t = 0, 1, \ldots, N-1$, and $N \in \mathbb{N}$, where $\tilde{r}_k > 0$, $\tilde{\varphi}_k \in [-\pi, \pi)$, and $\tilde{\omega}_k \in [-\pi, \pi)$ denote the amplitude, initial phase, and frequency, respectively, for the k:*th* signal component. Then, if there exists $\omega_0 \in [-\pi, \pi)$ such that $\tilde{\omega}_k = k\omega_0$, for $k = 1, \ldots, K$, the signal in (1) is referred to as being harmonic, with fundamental frequency, or pitch, ω_0. Note that this is the case also if some of the components are missing, i.e., $\tilde{r}_k = 0$ for some k, as the signal period is still $2\pi/\omega_0$. The model in (1) appear in many signal processing applications, not least in audio processing, and considerable effort has been directed towards deriving estimators for the parameters ω_0 and $(\tilde{r}_k, \tilde{\varphi}_k)$, for $k = 1, \ldots, K$ (see, e.g., [6] for an overview). However, for inharmonic signals, the integer relationship between the component frequencies is only approximate [8, 19]. That is,

$$\tilde{\omega}_k = \omega_0 k + \Delta_k, \;\; k = 1, \ldots, K, \tag{2}$$

where Δ_k are called inharmonicity parameters. As to distinguish from the case with completely unrelated sinusoidal components, it is here assumed that the inharmonicity parameters are small in the sense that $|\Delta_k| \ll \omega_0$. Thus, assuming that the component frequencies satisfy (2) constitutes a type of middle ground between the highly structured harmonic model, where $\Delta_k = 0$, for all k, and the unstructured sinusoidal model, where there is no relation between the frequency components. In some special cases, parametric models for the inharmonicity exists. For example, a common model for vibrating strings is

$$\tilde{\omega}_k = \omega_0 k \sqrt{1 + k^2\beta} \tag{3}$$

[1] In the interest of generality, we here consider complex-valued signals. For real-valued signals, a corresponding complex version may easily be formed as the discrete-time analytical signal [18].

where $\beta > 0$ is a parameter related to the stiffness of the string [7]. Thus, for this model, the inharmonicity parameters are given by $\Delta_k = \omega_0 k \left(\sqrt{1 + k^2\beta} - 1 \right) > 0$, for all k. However, such a structured model may not be assumed in the general case, and, in this work, we therefore do not assume any particular structure of $\{\Delta_k\}_{k=1}^K$, or indeed that any useful such structure exists. Rather, we aim to put the intuitive concept of inharmonic, i.e., almost harmonic, signals on more solid foundation by offering three conceptually different definitions of the meaning of a fundamental frequency for a non-periodic signal. Common for the three definitions will be the existence of a perfectly harmonic waveform,

$$\mu_t = \sum_{\ell=1}^{L} r_\ell e^{i\varphi_\ell + i\omega_0 \ell t}, \tag{4}$$

with L not necessarily being equal to K, that is a best approximation of x_t, with the definition of optimality differing between the definitions. In this framework, ω_0 will be the definition of the pitch for an inharmonic signal x_t. Furthermore, in order to find bounds on estimation performance, as well as to formulate estimators, we assume that the measured signal is well modeled as

$$y_t = x_t + e_t, \tag{5}$$

where e_t is a circularly symmetric white Gaussian noise with variance $\tilde{\sigma}^2$. However, most results presented herein may be readily extended to non-white noise processes. Throughout, let

$$\vartheta = \begin{bmatrix} \omega_0 & \varphi_1 & \dots & \varphi_L & r_1 & \dots & r_L \end{bmatrix}^T$$

denote the parameter vector defining the approximating signal, i.e., $\mu_t \equiv \mu_t(\vartheta)$, and let $\mathbf{y} = \begin{bmatrix} y_0 & \dots & y_{N-1} \end{bmatrix}^T$ be the vector of available signal samples.

3 ℓ_2 optimality and misspecified models

We initially consider approximating the inharmonic signal in (1) in the ℓ_2 sense, i.e.,

$$\vartheta_0 = \arg\min_{\vartheta} \frac{1}{N} \sum_{t=0}^{N-1} |x_t - \mu_t(\vartheta)|^2 . \tag{6}$$

That is, the approximating pitch is the harmonic waveform yielding the least squared deviation from the inharmonic signal. For notational brevity, $L = K$ for the ℓ_2 approximation.[2] Approximations such as (6) has earlier been applied in speech coding applications for decreasing the data rate in speech analysis/synthesis systems [21]. In addition to the intuitive appeal of this choice, as well as the tractability of computing ϑ, this approximation may be interpreted as a so-called pseudo-true parameter within the framework of misspecified models. Specifically, consider a scenario in which it is (erroneously) believed that the signal samples, y_t are perfectly harmonic, i.e., generated as $y_t = \mu_t + w_t$, where w_t is a circularly symmetric white Gaussian noise with (unknown) variance σ^2. Then, one may consider the following definition.

Definition 1 (Pseudo-true parameter [12])**.** Consider a signal sample $\mathbf{y}$. For a probability density function (pdf) p, parametrized by the parameter vector ϑ, the pseudo-true parameter, ϑ_0, is defined as

$$\vartheta_0 = \arg\min_{\vartheta} \; -\mathbb{E}_{\mathbf{y}} \left(\log p(\mathbf{y}; \vartheta) \right), \tag{7}$$

where $\mathbb{E}_{\mathbf{y}}$ denotes expectation with respect to the pdf of $\mathbf{y}$.

As may be noted, the pseudo-true parameter minimizes the Kullback-Leibler distance between the pdf of the assumed model and the actual pdf of the signal sample. In our case, the following proposition holds.

Proposition 1 (Pseudo-true parameter)**.** *Under the Gaussian assumption, the pseudo-true parameter ϑ_0 for the harmonic model is given by* (6)*, and the pseudo-true variance is*

$$\sigma^2 = \tilde{\sigma}^2 + \frac{1}{N} \sum_{t=0}^{N-1} |\xi_t(\vartheta_0)|^2 ,$$

where $\xi_t(\vartheta) \triangleq \mu_t(\vartheta) - x_t$.

Proof. As both the assumed and true distributions are Gaussian, the result follows directly. □

[2] This may be extended to $L \neq K$ in a straightforward manner, although $L = K$ allows for making additional comparisons to the perfectly harmonic case. It may be noted that $L < 2K$ is required for avoiding the so-called sub-octave problem [20].

From this, we may conclude that approximating the inharmonic signal in ℓ_2 may be interpreted as finding the Gaussian pdf with mean identical to a periodic waveform that best approximates the true signal pdf in the Kullback-Leibler sense. Furthermore, it can be shown that the misspecified MLE (MMLE), i.e., the MLE derived under an assumed model different from that of the actual measurements, asymptotically tends to the pseudo-true parameter, i.e., $\hat{\vartheta}_{MMLE} \to \vartheta_0$ almost surely as the signal to noise ratio (SNR), or number of signal samples, N, tends to infinity [12]. This leads to a very practical consequence: the harmonic waveform μ_t resulting from the ℓ_2 approximation in (6) corresponds to the expected result when applying an harmonic MLE, or approximations thereof [6], to inharmonic measurements. We summarize this in the following definition.

Definition 2. Let x_t, for $t = 0, \ldots, N-1$, be an inharmonic waveform. Then, the best harmonic approximation in ℓ_2 is given by $\mu_t(\vartheta_0)$, where ϑ_0 solves (6).

The harmonic signal in Def. 2 may be seen as the quantity being (tacitly) estimated when applying estimators derived under an harmonic assumption to inharmonic signals. Furthermore, by considering the interpretation as the pseudo-true parameter, one may find a bound on performance on any unbiased estimator of ϑ_0. Such a family of bounds is the misspecified CRLB (MCRLB) [13]. Specifically, considering estimators that satisfy the MLE unbiasedness conditions, the following theorem, adapted from [13], yields a bound on estimator variance.

Theorem 1. *Let $\hat{\vartheta}$ be an estimator of ϑ_0 that is unbiased under the signal pdf. Then,*

$$\mathbb{E}_{\mathbf{y}}\left((\hat{\vartheta}-\vartheta_0)(\hat{\vartheta}-\vartheta_0)^T\right) \succeq A(\vartheta_0)^{-1}F(\vartheta_0)A(\vartheta_0)^{-1} \tag{8}$$

where $A(\vartheta) = -\frac{\sigma^2}{\tilde{\sigma}^2}F(\vartheta) - \tilde{F}(\vartheta)$ and

$$F(\vartheta) = \frac{2\tilde{\sigma}^2}{(\sigma^2)^2}\sum_{t=0}^{N-1}\nabla_\vartheta\mu_t^{\mathfrak{R}}(\vartheta)\nabla_\vartheta\mu_t^{\mathfrak{R}}(\vartheta)^T + \nabla_\vartheta\mu_t^{\mathfrak{I}}(\vartheta)\nabla_\vartheta\mu_t^{\mathfrak{I}}(\vartheta)^T,$$

$$\tilde{F}(\vartheta) = \frac{2}{\sigma^2}\sum_{t=0}^{N-1}\left(\xi_t^{\mathfrak{R}}(\vartheta)\nabla^2_\vartheta\mu_t^{\mathfrak{R}}(\vartheta) + \xi_t^{\mathfrak{I}}(\vartheta)\nabla^2_\vartheta\mu_t^{\mathfrak{I}}(\vartheta)\right).$$

Here, $(\cdot)^{\mathfrak{R}} = \mathfrak{Re}(\cdot)$ and $(\cdot)^{\mathfrak{I}} = \mathfrak{Im}(\cdot)$ denote the real and imaginary parts, respectively.

Proof. See the appendix. □

The MCRLB is given by the diagonal of the right-hand side of (8) and thus provides a lower bound on the variance of any estimator of ϑ_0 that is unbiased under the pdf of the measured signal. It may be noted that when $\Delta_k = 0$, for all k, i.e., when the signal x_t is perfectly harmonic, $\tilde{F}(\vartheta_0) = 0$, and $F(\vartheta_0)$ is the standard Fisher information matrix (FIM). In this case, the MCRLB coincides with the CRLB of a harmonic signal. Thus, the MCRLB provides a means of assessing the performance of any estimator derived under the harmonic assumption, even when the observed signal is inharmonic. Further, as $N \to \infty$, one may express the MCRLB corresponding to the pseudo-true fundamental frequency in closed form, as detailed below.

Proposition 2 (Asymptotic MCRLB)**.** *Let the pseudo-true parameter be*

$$\vartheta_0 = \begin{bmatrix} \omega_0 & \varphi_1 & \dots & \varphi_K & r_1 & \dots & r_K \end{bmatrix}^T.$$

Then, as $N \to \infty$, the asymptotic MCRLB *for the pseudo-true fundamental frequency ω_0 is given by*

$$\mathrm{MCRLB}(\omega_0) = \tilde{\sigma}^2 \frac{C + E}{(C - E + Z + D)^2} \tag{9}$$

where $C = \frac{N(N^2-1)\sum_{k=1}^{K} k^2 r_k^2}{6}$, and

$$\begin{aligned}
Z &= -2\sum_{k=1}^{K} k^2 r_k^2 \frac{N(N-1)(2N-1)}{6} \\
&\quad + 2\sum_{k=1}^{K}\sum_{t=0}^{N-1} k^2 r_k \tilde{r}_k t^2 \cos(\breve{\varphi}_k + \breve{\omega}_k t) \\
D &= 2(N-1)\left[\frac{N(N-1)}{2}\sum_{k=1}^{K} k^2 r_k^2 - \sum_{k=1}^{K}\sum_{t=0}^{N-1} k^2 r_k \tilde{r}_k t \cos(\breve{\varphi}_k + \breve{\omega}_k t)\right] \\
E &= \frac{2}{N}\sum_{k=1}^{K} k^2 \tilde{r}_k^2 \left(\sum_{t=0}^{N-1} t \sin(\breve{\varphi}_k + \breve{\omega}_k t)\right)^2 \\
&\quad + \frac{2}{N}\sum_{k=1}^{K} k^2 \left(\tilde{r}_k \sum_{t=0}^{N-1} t \cos(\breve{\varphi}_k + \breve{\omega}_k t) - r_k \frac{N(N-1)}{2}\right)^2,
\end{aligned}$$

where $\breve{\varphi}_k = \varphi_k - \tilde{\varphi}_k$ *and* $\breve{\omega}_k = k\omega_0 - \tilde{\omega}_k$, *for* $k = 1, \ldots, K$.

Proof. See the appendix. □

It may be noted that in the perfectly harmonic case, the terms E, Z, and D are all equal to zero as the pseudo-true parameter ϑ_0 then coincides with the actual signal parameter. It is worth noting that $\tilde{\sigma}^2/C$ is equal to the asymptotic CRLB for the perfectly harmonic model [22].

Remark 1. As one may compute the MCRLB for an estimate of the pseudo-true fundamental frequency ω_0, it is also possible to construct a bound for the expected MSE for misspecified estimators. That is, if one considers an estimate of ω_0 to be a misspecified estimate of $\tilde{\omega}_1$, the theoretical MSE is given by

$$\mathbb{E}_{\mathbf{y}}\left((\hat{\omega}_0 - \tilde{\omega}_1)^2\right) = \mathrm{MCRLB}(\omega_0) + (\omega_0 - \tilde{\omega}_1)^2. \tag{10}$$

In fact, as will be illustrated in the numerical section, the MSE for this misspecified estimate may for moderate values of SNR and sample lengths N be lower than the CRLB for an unstructured sinusoidal model, due to the MCRLB often being considerable smaller than the CRLB.

Despite its appeal as formalizing tacit assumptions behind using harmonic estimators, as well as resulting in easily computed performance bounds, Def. 2 becomes unsatisfactory if one allows for very long signals. Although the normalization by N allows for considering limits of the criterion (6), as it is guaranteed to be finite, the definition of pitch becomes ambiguous. To see this, it may be noted that the correlation of any two sinusoids with distinct frequencies tends to zero as $N \to \infty$. Thus, (6) is minimized by setting a multiple of the pseudo-true ω_0 equal to the frequency $\tilde{\omega}_k$ corresponding to the largest amplitude $\tilde{r}_k$. However, any integer multiple is an equally valid choice. Def. 2 allows for selecting ω_0 as $\omega_0 = \tilde{\omega}_k/\ell$, where $k = \arg\max_m \tilde{r}_m$, for any $\ell \in \{1, \ldots, K\}$. Thus, K indistinguishable candidates for ω_0 exist if N is large enough. It may also be noted that ω_0 depends on the initial phases $\tilde{\varphi}_k$ of the inharmonic signal, parameters that may be considered nuisance.

The issue of ambiguity for large N is illustrated in Figures 1 and 2. Specifically, Figure 1 displays the ℓ_2 cost function for a perfectly harmonic signal with $\omega_0 = \pi/10$ consisting of five harmonics, where $\tilde{r}_2 = \tilde{r}_3$ are the largest amplitudes, with the top and bottom panels of Figure 1 showing the cost for $N = 500$ and $N = 3000$, respectively. As can be seen, for both cases, ω_0 corresponds to

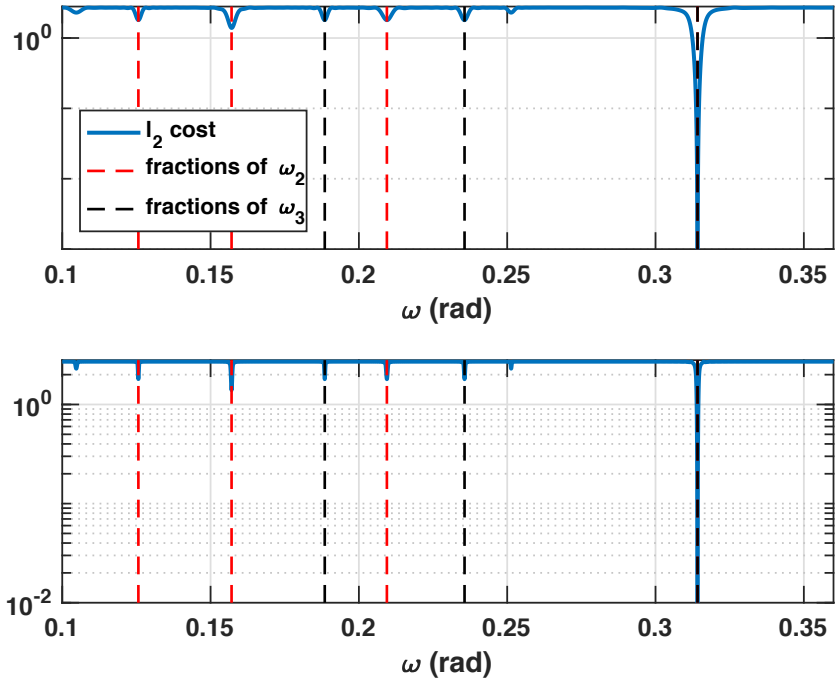

Figure 1: The ℓ_2 cost function used in Def. 2 when applied to a perfectly harmonic signal with fundamental frequency $\omega_0 = \pi/10$ for two different sample lengths. Top panel: $N = 500$. Bottom panel: $N = 3000$.

the unique global minimum, with a cost of exactly zero. In contrast, Figure 2 displays the same scenario, with the difference being that the perfectly harmonic structure having been replaced with the string model in (3) with $\beta = 10^{-3}$. For $N = 500$, the global minimum is still unique, and the definition thus unambiguous. However, for $N = 3000$, the cost function value at the local minimum at $\tilde{\omega}_2/4$ approaches that of the global minimum. If one lets $N \to \infty$, these cost functions values will become identical, in addition to several other isolated global minima appearing. The issue of ambiguity may be addressed by instead considering the spectral properties of the signal x_t, allowing for defining a harmonic approximation of the more abstract signal, i.e., when no particular sample length N has been specified. This can be achieved through the use of OMT, as described next.

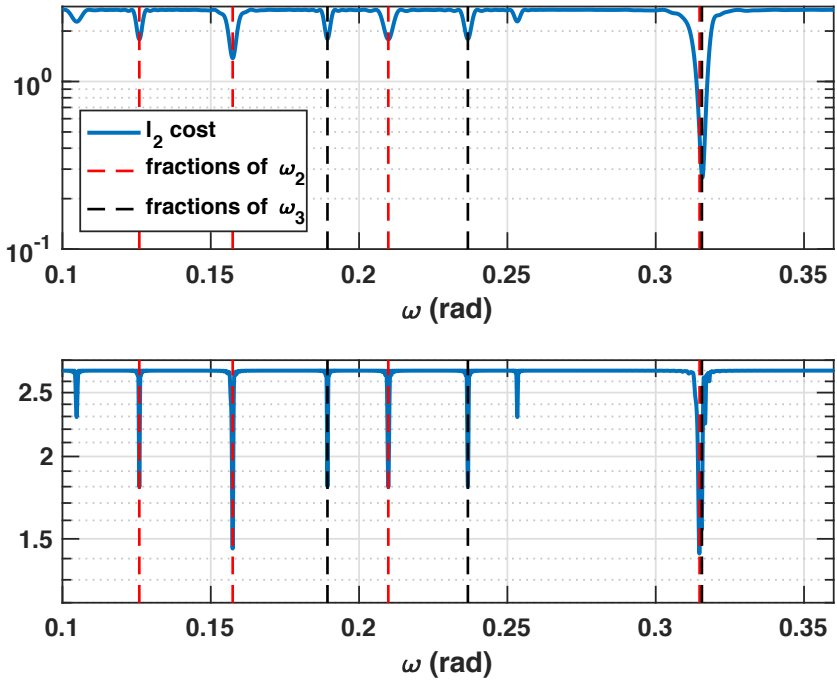

Figure 2: The ℓ_2 cost function used in Def. 2 when applied to an inharmonic signal generated using (3), with $\omega_0 = \pi/10$ and $\beta = 10^{-3}$ for two different sample lengths. Top panel: $N = 500$. Bottom panel: $N = 3000$.

4 An OMT-based definition of pitch

As an alternative to defining the harmonic counterpart of an inharmonic signals by means of waveform approximation, such as in Def. 2, one may instead consider the signals' spectral representation. Specifically, the spectrum of the signal in (1) is given by

$$\Phi_x(\omega) = 2\pi \sum_{k=1}^{K} \tilde{r}_k^2 \delta(\omega - \tilde{\omega}_k), \tag{11}$$

where $\omega \in [-\pi, \pi)$, and $\delta(\cdot)$ denotes the Dirac delta function. This constitutes a more abstract representation of x_t as it is not related to any particular sample

length N. Furthermore, any harmonic spectrum may be represented as

$$\Phi_\mu(\omega) = 2\pi \sum_{\ell=1}^{L} r_\ell^2 \delta(\omega - \ell\omega_0), \tag{12}$$

where we let L be potentially different from K. Thus, one may define the fundamental frequency of an inharmonic signal as the ω_0 corresponding to an approximating harmonic spectrum Φ_μ. Note however that considering the approximation in L_p, for some $p \geq 1$, leads to same problem as was encountered in the limiting case for ℓ_2 approximation of the waveform: selecting ω_0 as $\omega_0 = \tilde{\omega}_k/\ell$, where $k = \arg\max_m \tilde{r}_m$, for any $\ell \in \{1, \ldots, L\}$ yield equally good approximations. However, a measure of distance not related to point-wise comparison of Φ_x and Φ_μ may be obtained by considering the framework of OMT. Within this framework, the best approximation corresponds to the one requiring the least costly perturbation in order to shift the observed signal to that of a perfectly harmonic model, with perturbations being realized by moving the whole distribution, as opposed to the point-wise changes implied by L_p norms.

To formulate this, let $\mathbb{T} = [-\pi, \pi)$ and let $\mathcal{M}_+(\mathbb{T})$ denote the set of non-negative, generalized integrable functions on $\mathbb{T}$. Elements of $\mathcal{M}_+(\mathbb{T})$ may be interpreted as distributions of mass on $\mathbb{T}$, and, in particular, $\Phi_x \in \mathcal{M}_+(\mathbb{T})$ and $\Phi_\mu \in \mathcal{M}_+(\mathbb{T})$. Then, for any $\Phi_0, \Phi_1 \in \mathcal{M}_+(\mathbb{T})$ with the same total mass, i.e., $\int_\mathbb{T} \Phi_0(\omega)d\omega = \int_\mathbb{T} \Phi_1(\omega)d\omega$, one may define a notion of distance, $S : \mathcal{M}_+(\mathbb{T}) \times \mathcal{M}_+(\mathbb{T}) \to \mathbb{R}$, between them by the Monge-Kantorovich problem of OMT [14],

$$\begin{aligned} S(\Phi_0, \Phi_1) = \min_{\mathcal{M}_+(\mathbb{T}\times\mathbb{T})} & \int_{\mathbb{T}\times\mathbb{T}} M(\omega_1, \omega_2) c(\omega_1, \omega_2) d\omega_1 d\omega_2 \\ \text{s.t.} & \int_\mathbb{T} M(\cdot, \omega)d\omega = \Phi_0 \ , \ \int_\mathbb{T} M(\omega, \cdot)d\omega = \Phi_1, \end{aligned}$$

where $c : \mathbb{T} \times \mathbb{T} \to \mathbb{R}_+$ is a cost function defining the cost of moving a unit mass. Here, M is referred to as a transport plan as it may be interpreted as describing how mass is moved from Φ_0 to Φ_1. As the constraints ensure that M transports all available mass, and no other, between Φ_0 and Φ_1, the objective corresponds to the total cost of transport. The idea of using OMT as a measure of distance between spectra has earlier been considered in [23], wherein it was shown that S, for certain choices of c, may be used for defining a metric on $\mathcal{M}_+(\mathbb{T})$. Herein, we will use $c(\omega_1, \omega_2) = (\omega_1 - \omega_2)^2$, i.e., the cost of transport between two

frequencies is equal to their squared Euclidean distance. With this, one may find the best harmonic approximation of Φ_x in the OMT sense by minimizing $S(\cdot, \Phi_x)$ over the set of harmonic spectra. Note that, in contrast to L_p minimization, this corresponds to finding the most efficient way of perturbing the spectral peaks of Φ_x in frequency so that the resulting spectrum is harmonic, with the cost of moving a peak being proportional to its power. Formally,

$$\Phi_\mu = \arg\min_{\Phi \in \Omega_L} S(\Phi, \Phi_x), \tag{13}$$

where Ω_L is the set of harmonic spectra, i.e.,

$$\Omega_L = \Big\{ \Phi \in \mathcal{M}_+(\mathbb{T}) \mid \Phi(x) = 2\pi \sum_{\ell=1}^{L} r_\ell^2 \delta(\omega - \ell\omega_0) \ , \ r_\ell \geq 0, \omega_0 \in \mathbb{T} \Big\}.$$

To see that (13) may be solved efficiently, note that for any candidate ω_0, corresponding to a subset of Ω_L, all power at the frequency $\tilde{\omega}_k$ in Φ_x will be transported to the nearest integer multiple of ω_0 when evaluating S. That is, letting Ω_{L,ω_0} denote such a subset, i.e., all harmonic spectra with fundamental frequency ω_0, we get

$$\min_{\Phi \in \Omega_{L,\omega_0}} S(\Phi, \Phi_x) = 2\pi \sum_{k=1}^{K} \tilde{r}_k^2 \min_{\ell \in \{1,2,\ldots,L\}} (\ell\omega_0 - \tilde{\omega}_k)^2.$$

Thus, solving (13) is equivalent to solving

$$\operatorname*{minimize}_{\omega_0} \ 2\pi \sum_{k=1}^{K} \tilde{r}_k^2 \min_{\ell \in \{1,2,\ldots,L\}} (\ell\omega_0 - \tilde{\omega}_k)^2. \tag{14}$$

Furthermore, at least one harmonic spectrum attaining the minimal cost exists, and is given by

$$\Phi_\mu(\omega) = 2\pi \sum_{\ell=1}^{L} \left(\sum_{k \in \mathcal{I}_\ell} \tilde{r}_k^2 \right) \delta(\omega - \ell\omega_0),$$

where ω_0 solves (14), and with

$$\mathcal{I}_\ell = \left\{ k \mid \ell = \arg\min_m (m\omega_0 - \tilde{\omega}_k)^2 \right\} \tag{15}$$

denoting the set of indices that are transported to harmonic ℓ, for $\ell = 1, \ldots, L$. As may be noted, the maximal harmonic order, L, may not be equal to K (indeed, K may be unknown, or may not be a suitable choice for L). Although L could be left as a user defined parameter, we here offer a data dependent choice with both intuitive appeal and beneficial practical consequences.

Definition 3 (Maximal harmonic order)**.** For a set $\{\tilde{\omega}_k\}_{k=1}^K$ such that $\tilde{\omega}_{k+1} > \tilde{\omega}_k$ for $k = 1, \ldots, K-1$, let $d = \min\{\tilde{\omega}_1, \tilde{\omega}_2 - \tilde{\omega}_1, \tilde{\omega}_3 - \tilde{\omega}_2, \ldots, \tilde{\omega}_K - \tilde{\omega}_{K-1}\}$ i.e., the minimum distance between two consecutive frequencies. Then, the harmonic order for the set $\{\tilde{\omega}_k\}_{k=1}^K$ is defined as $L = \min\{\ell \in \mathbb{N} \mid \ell d \geq \tilde{\omega}_K\}$.

With this, the definition of the OMT harmonic spectrum is:

Definition 4 (Closest harmonic spectrum)**.** Let $\tilde{\Phi}(x) = 2\pi \sum_{k=1}^K \tilde{r}_k^2 \delta(x - \tilde{\omega}_k)$ be a (possibly) inharmonic spectrum, and let L be the maximal harmonic order as defined in Def. 3. Then, the closest harmonic spectrum (CHS) is defined as

$$\Phi_\mu(\omega) = 2\pi \sum_{\ell=1}^{L} \left(\sum_{k \in \mathcal{I}_\ell} \tilde{r}_k^2 \right) \delta(\omega - \ell\omega_0),$$

where ω_0 solves (14) and $\mathcal{I}_\ell$ is given by (15).

It may here be noted that selecting L according to Def. 3 acts as a safeguard against the so called sub-octave problem, i.e., associating the spectral lines with $\omega_0/2^P$, for $P \geq 1$, which is bound to happen if L is chosen excessively large. The following two propositions verify that Def. 4 behaves in a stable and predictable way. To simplify the proofs, define

$$q_L(\omega_0) \triangleq 2\pi \sum_{k=1}^{K} \tilde{r}_k^2 \min_{\ell \in \{1,2,\ldots,L\}} (\ell\omega_0 - \tilde{\omega}_k)^2, \tag{16}$$

i.e., the minimal cost of transporting Φ_x to a harmonic spectrum with fundamental frequency ω_0 and L harmonics. Thus, the fundamental frequency of the CHS minimizes q_L. Furthermore, it may be noted that the following propositions hold also if one instead of using Def. 3 selects $L = K$.

Proposition 3. *Let $\Phi_x(\omega) = 2\pi \sum_{k=1}^{K} \tilde{r}_k^2 \delta(\omega - k\tilde{\omega}_0)$ be a harmonic spectrum with fundamental frequency $\tilde{\omega}_0$. Then, the CHS is identical to Φ_x, and, in particular, $\omega_0 = \tilde{\omega}_0$.*

Proof. Clearly, the maximal harmonic order is $L = K$. Further, $q_K(\tilde{\omega}_0) = 0$, and $q_K(\omega') > 0$ for any $\omega' \neq \tilde{\omega}_0$. □

This proposition ensures that for a perfectly harmonic spectrum, the CHS is the spectrum itself. As may be seen from (16), the function is q_L is non-convex, with several local minima. Thus, for arbitrarily large inharmonicity parameters Δ_k and arbitrary choices of amplitudes $\tilde{r}_k$, the fundamental frequency cannot be guaranteed to be found in a certain region (it may be noted that this is also the case for the ℓ_2 approximation). However, for small harmonic perturbations, the following proposition holds.

Proposition 4. *Let $\tilde{\Phi}(\omega) = 2\pi \sum_{k=1}^{K} \tilde{r}_k^2 \delta(\omega - \tilde{\omega}_k)$, where $\tilde{\omega}_k = k\tilde{\omega}_0 + \Delta_k$, with $\|\Delta\|_\infty \triangleq \max_k |\Delta|_k < \tilde{\omega}_0/(2K+3)$. Then, $L \in \{K, K+1\}$, and the transport cost q_L has a exactly one local minimum ω_0 on the interval $\left(\tilde{\omega}_0 - \|\Delta\|_\infty, \tilde{\omega}_0 + \|\Delta\|_\infty\right)$. Furthermore,*

$$|\omega_0 - \tilde{\omega}_0| \leq \frac{\sum_{k=1}^{K} \tilde{r}_k^2 k}{\sum_{k=1}^{K} \tilde{r}_k^2 k^2} \|\Delta\|_\infty .$$

Proof. See the appendix. □

For non-extreme choices[3] of the amplitudes $\tilde{r}_k$, the local minimum in Prop. 4 is expected to be also the global minimum, i.e., small inharmonic perturbations is expected to yield a CHS whose pitch is close to that of the unperturbed perfectly harmonic signal. With this, one may conclude that Def. 4 provides a reasonable and well-behaved definition of pitch for inharmonic signals with some advantages over the ℓ_2 approximation. Firstly, as the approximation is performed in the spectral domain, there is no dependency parameters related to the observation of an instance of the signal, such as sample length or initial phases. Secondly, it codifies the intuitive idea that a signal is almost harmonic if only a slight perturbation of its spectrum is needed to obtain an harmonic structure.

[3] Adversarial examples with combinations of very large and small amplitudes can always be constructed as to move the global mimimum.

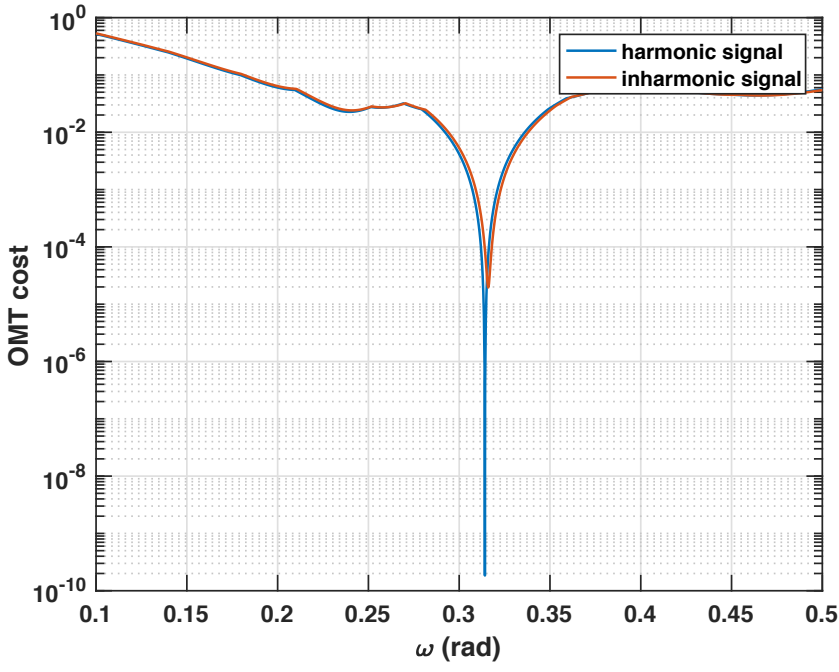

Figure 3: The OMT cost function, q_L, for a perfectly harmonic signal with fundamental frequency $\omega_0 = \pi/10$, as well as for an inharmonic signal generated from the string model in (3) with $\omega_0 = \pi/10$ and $\beta = 10^{-3}$.

To illustrate the intuitive appeal of Def. 4, Figure 3 considers the same example as in Figures 1 and 2, with the difference being that the approximation is performed in the spectral domain and in the OMT sense instead of in the temporal domain and ℓ_2. As can be seen, using Def. 4 causes only a small perturbation of the ω_0 defining the approximating spectrum as compared to the perfectly harmonic case. In fact, the harmonic spectrum of Def. 4 is constructed by slightly shifting the frequency locations of the peaks of the actual signal spectrum. This is illustrated in Figure 4, displaying the spectrum[4] of the inharmonic signal, as well as the spectra of the approximations from Defs. 2 and 4. As may be noted, the total power is retained for Def. 4, whereas amplitudes are underestimated or components altogether missing due to orthogonality for Def. 2. As a side, it

[4]As the spectra are all singular, the pointmasses are here represented by scaled arrows.

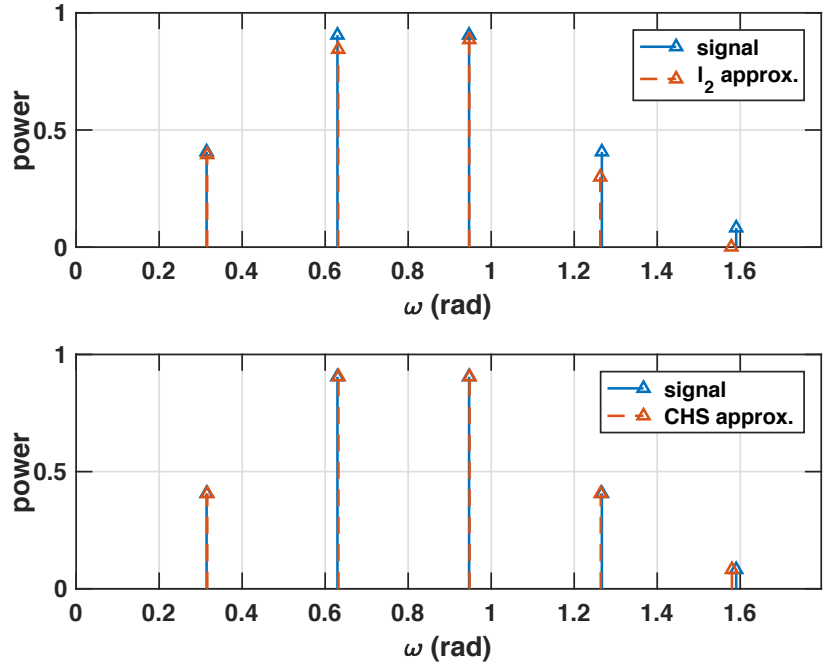

Figure 4: Spectrum of inharmonic signal, as well as spectra of ℓ_2 and CHS approximations corresponding to Defs. 2 and 4, respectively. Top panel: signal and ℓ_2 approximation. Bottom panel: signal and CHS approximation.

may also be noted that the cost function is quadratic in a neighborhood of this fundamental frequency.

As the ℓ_2 approximation is closely related to the MLE derived under the perfectly harmonic assumption, Def. 2 is readily applicable to actual estimation problems. However, Def. 4 may also be used as a plug-in estimator by replacing the quantities $(\tilde{r}_k, \tilde{\omega}_k)$ by finite-sample estimates $(\hat{\tilde{r}}_k, \hat{\tilde{\omega}}_k)$ and consider the obtained fundamental frequency to be an estimate of ω_0. The following proposition details the asymptotic behavior of such an estimator.

Proposition 5 (CHS estimate)**.** *Let $(\hat{\tilde{r}}_k, \hat{\tilde{\omega}}_k)$ be MLEs of $(\tilde{r}_k, \tilde{\omega}_k)$, for all k, obtained from an unstructured sinusoidal model with Gaussian noise. Then, under the assumptions of Prop 4, the plug-in estimate of the CHS pitch, $\hat{\omega}_0$, is an asymptotically, i.e.,*

as $N \to \infty$, consistent estimator of the CHS ω_0 with asymptotic variance given by

$$\mathrm{Var}(\hat{\omega}_0) = \frac{6\tilde{\sigma}^2}{N(N^2-1)\sum_{k=1}^{K} k^2\tilde{r}_k^2} + \frac{2\tilde{\sigma}^2}{N\left(\sum_{k=1}^{K} k^2\tilde{r}_k^2\right)^4}\sum_{k=1}^{K} k^2\tilde{r}_k^2\left(\sum_{\ell=1}^{K}\ell\tilde{r}_\ell^2(\ell\tilde{\omega}_k - k\tilde{\omega}_\ell)\right)^2. \tag{17}$$

Proof. See the appendix. □

Remark 2. It may be noted that the first term of (17) is identical to the asymptotic CRLB for the pitch in a perfectly harmonic model [22]. Further, in the perfectly harmonic case, i.e., $\tilde{\omega}_k = k\omega_0$, for all k, the second term of (17) is equal to zero, i.e., the CHS plug-in estimate has the same asymptotic performance as the MLE for the perfectly harmonic model. This is not surprising; in the perfectly harmonic case and under the assumption of additive Gaussian white noise, the criterion q_L minimized by the CHS is asymptotically equivalent to the EXIP cost function corresponding to finding an estimate of the pitch from a set of unstructured frequency estimates [17]. Furthermore, the EXIP estimate has been shown to asymptotically have the same performance as the MLE [15].

5 Stochastic representation

As may be noted, the parameters of the inharmonic signal in (1) have in Defs. 2 and 4 been assumed to be deterministic (although the spectral representation allows for, e.g., the initial phases to be random), and, as a consequence, the corresponding harmonic structure has been defined in terms of approximations. In the third and last definition presented herein, we take an alternative approach and view (1) as a realization of a stochastic process in which the frequency parameters are random. Specifically, for the frequencies $\tilde{\omega}_k$ in (2), we let ω_0 be a deterministic parameter, whereas Δ_k are independent zero-mean random variables so that

$$\mathbb{E}(\tilde{\omega}_k) = \mathbb{E}(k\omega_0 + \Delta_k) = k\omega_0, \text{ for } k = 1, \ldots, K.$$

With this, ω_0 may be interpreted as the fundamental frequency in an average sense[5]. In order to arrive at a model allowing for manipulations, we will herein

[5]Note, however, that the expectation of x_t is not a periodic waveform.

further assume that the inharmonicity parameters may be well modeled as Gaussian random variables, i.e., $\Delta_k \in \mathcal{N}(0, \sigma_\Delta^2)$, where σ_Δ^2 denotes the common variance[6]. It may be noted that this representation constitutes a weaker, less structured assumption on the nature of the inharmonic deviations than that provided of, e.g., the stiff string model in (3), and may therefore be used to model inharmonic signals for which there is no known parametric description. We summarize this in the following definition.

Definition 5 (Expected harmonic signal)**.** The signal in (1) is a realization of the random variable

$$x_t = \sum_{k=1}^{K} \tilde{r}_k e^{i\tilde{\varphi}_k + i(k\omega_0 + \Delta_k)t}, \tag{18}$$

where the pdf of $\Delta = \begin{bmatrix} \Delta_1 & \dots & \Delta_K \end{bmatrix}^T$ is given by

$$p(\Delta) = \frac{1}{(2\pi\sigma_\Delta^2)^{K/2}} \exp\left(-\frac{1}{2\sigma_\Delta^2} \|\Delta\|_2^2\right).$$

The corresponding harmonic signal is given by (4).

It may be noted that there is an interesting relation between Defs. 4 and 5. Specifically, for a given realization of x_t, corresponding to an observation of $\tilde{\omega}_k$, the waveform may be explained perfectly by matching amplitudes, initial phases, and by making any choice of ω_0 and Δ such that

$$k\omega_0 + \Delta_k = \tilde{\omega}_k, \text{ for } k = 1, \dots, K.$$

However, the choice that is most likely, in the sense of maximizing $p(\Delta)$, is given by

$$\omega_0 = \arg\min_{\omega} \sum_{k=1}^{K} (k\omega - \tilde{\omega}_k)^2 = \frac{\sum_{k=1}^{K} k\tilde{\omega}_k}{\sum_{k=1}^{K} k^2},$$

and $\Delta_k = \tilde{\omega}_k - \omega_0/k$, for all k. That is, the estimate of μ_t corresponds to a variation of Def. 4 in which the cost of transport is not related to the component

[6]This choice is made to simplify the exposition. However, all results herein may be readily extended to allowing the variance to differ among the inharmonicity parameters.

power, i.e., as if $\tilde{r}_k = 1$, for all k. It may also be noted that this model is identifiable for all finite σ_Δ^2, i.e., one is not forced to identify one of the frequencies $\tilde{\omega}_k$ with an integer multiple of ω_0 in order for the model to be well-defined.

Considering the noisy observation model in (5), we may proceed by asking what type of performance bounds and estimators that are relevant for the signal in Def. 5. To this end, let $\breve{\vartheta} = \begin{bmatrix} \vartheta^T & \Delta^T \end{bmatrix}^T$ denote the concatenation of the deterministic vector ϑ and the stochastic vector Δ. Furthermore, let $\mathbf{x}(\breve{\vartheta})$ denote the vector consisting of the N signals samples in (18), parametrized by $\breve{\vartheta}$. Then, assuming that the measurement noise is independent of Δ, the joint pdf of the measurement $\mathbf{y}$ and the inharmonicity Δ is given by

$$p(\mathbf{y}, \Delta; \vartheta) = p(\mathbf{y} \mid \Delta; \vartheta) p(\Delta), \tag{19}$$

where

$$p(\mathbf{y} \mid \Delta; \vartheta) = \frac{1}{(\pi\tilde{\sigma}^2)^N} \exp\left(-\frac{1}{\tilde{\sigma}^2} \|\mathbf{y} - \mathbf{x}(\vartheta, \Delta)\|_2^2 \right)$$

is the conditional pdf of the measurement. From this, it may be noted that it is not trivial to compute the CRLB for ϑ, nor to derive the MLE, as the marginal density $p(\mathbf{y}; \vartheta)$ is not available; this requires computing a K dimensional integral with a non-linear integrand. Also, the Bayesian CRLB is not applicable in this case as prior distributions are only available for a subset of the parameters. However, it is possible to find a performance bound for the model in (19) by means of the HCRLB [24, 25]. Furthermore, this bound may, as we will see in the numerical section, be asymptotically attained by a hybrid ML/MAP estimator [26]. The following result adapted from [24] holds.

Proposition 6 (Hybrid Cramér-Rao lower bound [24])**.** *Let $\hat{\breve{\vartheta}}$ be an unbiased estimator of $\breve{\vartheta}$ in the sense that $\mathbb{E}_{\mathbf{y},\Delta}(\hat{\vartheta}) = \vartheta$ and $\mathbb{E}_{\mathbf{y},\Delta}(\hat{\Delta}) = \mathbb{E}_\Delta(\Delta)$, for any ϑ. Then,*

$$\mathbb{E}_{\mathbf{y},\Delta}\left((\hat{\breve{\vartheta}} - \breve{\vartheta})(\hat{\breve{\vartheta}} - \breve{\vartheta})^T \right) \succeq \breve{F}(\breve{\vartheta})^{-1}, \tag{20}$$

where

$$\breve{F} = \mathbb{E}_{\mathbf{y},\Delta}\left(\nabla_{\breve{\vartheta}} \log p(\mathbf{y}, \Delta; \vartheta) \, \nabla_{\breve{\vartheta}} \log p(\mathbf{y}, \Delta; \vartheta)^T \right). \tag{21}$$

Here, $\mathbb{E}_{\mathbf{y},\Delta}$ and $\mathbb{E}_\Delta$ denote expectation with respect to the joint pdf of $\mathbf{y}$ and Δ, and the marginal pdf of Δ, respectively.

For non-linear measurement models, such as the one considered herein, the HCRLB is in general only tight asymptotically [25]. However, it can be shown that if the bound is tight, then it is attained by the hybrid ML/MAP estimator [25]. This property makes the HCRLB attractive, especially for the case considered herein, where it may be computed easily, as detailed in the following proposition. Furthermore, the ML/MAP allows for straightforward implementation.

Proposition 7. *For the inharmonic pitch model in* (19)*, the matrix* $\breve{F}$ *in* (21) *defining the HCRLB is given by*

$$\breve{F} = \mathbb{E}_{\Delta}\left(F(\breve{\vartheta})\right) + \begin{bmatrix} 0 & 0 \\ 0 & \frac{1}{\sigma_{\Delta}^2} I \end{bmatrix}, \tag{22}$$

where

$$F(\breve{\vartheta}) = \frac{2}{\tilde{\sigma}^2} \sum_{t=0}^{N-1} \nabla_{\breve{\vartheta}} \mu_t^{\Re}(\breve{\vartheta}) \nabla_{\breve{\vartheta}} \mu_t^{\Re}(\breve{\vartheta})^T + \nabla_{\breve{\vartheta}} \mu_t^{\Im}(\breve{\vartheta}) \nabla_{\breve{\vartheta}} \mu_t^{\Im}(\breve{\vartheta})^T,$$

and I *denotes the identity matrix of size* $K \times K$*. A detailed expression for* $\mathbb{E}_{\Delta}(F(\breve{\vartheta}))$ *may be found in the appendix.*

Proof. See the appendix. □

Remark 3. Partitioning $\breve{F}$ as

$$\breve{F} = \begin{bmatrix} F_{\vartheta,\vartheta} & F_{\vartheta,\Delta}^T \\ F_{\vartheta,\Delta} & F_{\Delta,\Delta}, \end{bmatrix}$$

where $F_{\vartheta,\vartheta}$ is the $(2K+1) \times (2K+1)$ block corresponding to the non-random parameters and $F_{\Delta,\Delta}$ corresponds to the K inharmonicity parameters, it may be noted $F_{\vartheta,\vartheta}$ converges to the FIM corresponding to a perfectly harmonic model when letting $\sigma_{\Delta}^2 \to 0$. Also, $\frac{1}{N} F_{\Delta,\Delta} \approx \frac{1}{N} \frac{1}{\sigma_{\Delta}^2} I$ when $\sigma_{\Delta}^2 \to 0$ and N is reasonably large. Noting that the Schur complement of $F_{\Delta,\Delta}$ in $\breve{F}$ is

$$F_{\vartheta,\vartheta} - F_{\vartheta,\Delta}^T F_{\Delta,\Delta}^{-1} F_{\vartheta,\Delta}$$

and noting that $F_{\vartheta,\Delta}$ converges to a finite matrix as $\sigma_{\Delta}^2 \to 0$, the top $(2K+1) \times (2K+1)$ block of $\breve{F}^{-1}$ converges to $F_{\vartheta,\vartheta}^{-1}$, as

$$F_{\vartheta,\Delta}^T F_{\Delta,\Delta}^{-1} F_{\vartheta,\Delta} \to \sigma_{\Delta}^2 F_{\vartheta,\Delta}^T F_{\vartheta,\Delta} \to 0$$

when $\sigma_\Delta^2 \to 0$. That is, the HCRLB converges to the CRLB of the perfectly harmonic model as $\sigma_\Delta^2 \to 0$.

Furthermore, it can be shown that, given some regularity conditions on the pdf p [25], the HCRLB is asymptotically tight and asymptotically attained by the hybrid ML/MAP estimator, i.e., by $\arg\max_{\vartheta,\Delta} p(\mathbf{y}, \Delta; \vartheta)$. Thus, the bound in Prop. 7 constitutes a useful predictor of estimation performance. Next, we present the ML/MAP estimator for the inharmonic pitch model and show that it lends itself to straightforward implementation. From (19), it may be noted the ML/MAP estimate of $(\vartheta, \tilde{\sigma}^2, \Delta)$ maximizes the function

$$\mathcal{L} = -N \log \tilde{\sigma}^2 - \frac{1}{\tilde{\sigma}^2} \left\| \mathbf{y} - \mathbf{x}(\breve{\vartheta}) \right\|_2^2 - \frac{1}{2\sigma_\Delta^2} \left\| \Delta \right\|_2^2,$$

which is the log-likelihood of (19), excluding constant terms. It may be noted that we here are required to estimate also the noise variance $\tilde{\sigma}^2$. In order to formulate the ML/MAP estimator, define the dictionary function $A : \mathbb{R}^K \to \mathbb{C}^{N \times K}$

$$A(\boldsymbol{\omega}) = \begin{bmatrix} a(\omega_1) & \dots & a(\omega_K) \end{bmatrix},$$

where $\boldsymbol{\omega} = \begin{bmatrix} \omega_1 & \dots & \omega_K \end{bmatrix}^T$, and $a : \mathbb{R} \to \mathbb{C}^N$ is the Fourier vector.

Proposition 8 (ML/MAP estimator)**.** *Let* $\boldsymbol{\omega}$ *be the set of frequencies maximizing the function*

$$\psi_{\text{ML/MAP}}(\boldsymbol{\omega}) = -N \log \Sigma(\boldsymbol{\omega}) - \frac{1}{2\sigma_\Delta^2} \nu(\boldsymbol{\omega}),$$

where

$$\Sigma(\boldsymbol{\omega}) = \frac{1}{N} \left\| \mathbf{y} - A(\boldsymbol{\omega}) \left(A(\boldsymbol{\omega})^H A(\boldsymbol{\omega}) \right)^{-1} A(\boldsymbol{\omega})^H \mathbf{y} \right\|_2^2,$$
$$\nu(\boldsymbol{\omega}) = \sum_{k=1}^{K} \left(\omega_k - k \frac{\sum_{\ell=1}^{K} \ell \omega_\ell}{\sum_{\ell=1}^{K} \ell^2} \right)^2.$$

Then, the ML/MAP estimates of the fundamental frequency and inharmonicity parameters are given by

$$\omega_0 = \frac{\sum_{k=1}^{K} k \omega_k}{\sum_{k=1}^{K} k^2}$$

and $\Delta_\ell = \omega_\ell - \ell\omega_0$*, for* $\ell = 1, \ldots, K$*, respectively. Furthermore, the estimates of the noise variance and the vector of complex amplitudes are given by* $\tilde{\sigma}^2 = \Sigma(\boldsymbol{\omega})$ *and*

$$\left(A(\boldsymbol{\omega})^H A(\boldsymbol{\omega})\right)^{-1} A(\boldsymbol{\omega})^H \mathbf{y},$$

respectively.

Proof. See appendix. □

Remark 4. It may be noted that the ML/MAP estimate of the model parameters are found by maximizing $\psi_{\text{ML/MAP}}$ over a set of K unconstrained frequencies, which may be realized by a non-linear search. As $\psi_{\text{ML/MAP}}$ is non-convex, such a search requires a good initial point. In favorable noise conditions, such an initial point may be obtained by simple peak-picking in the periodogram.

As can be seen from the criterion $\psi_{\text{ML/MAP}}$, one may obtain two extreme cases by letting $\sigma_\Delta^2 \to \infty$ and $\sigma_\Delta^2 \to 0$, respectively. In the former case, maximizing $\psi_{\text{ML/MAP}}$ becomes equivalent to minimizing Σ with respect to a set of unconstrained frequencies $\boldsymbol{\omega}$, i.e., one obtains the MLE for a model with K unrelated sinusoids. In the latter case, one in the limit arrives at the problem

$$\underset{\boldsymbol{\omega}}{\text{maximize}} \; -\Sigma(\boldsymbol{\omega}) \;, \text{ s.t. } \nu(\boldsymbol{\omega}) = 0,$$

or, equivalently,

$$\underset{\omega_0, \boldsymbol{\omega}}{\text{maximize}} \; -\Sigma(\boldsymbol{\omega}) \;, \text{ s.t. } \omega_k = k\omega_0, \text{ for } k = 1, \ldots, K,$$

i.e., the misspecified MLE corresponding to a perfectly harmonic approximation.

6 Numerical examples

In this section, we provide numerical examples illustrating the derived theoretical results.

6.1 Deterministic waveform

To illustrate the behavior of Def. 2 and the implied MMLE for varying degrees of inharmonicity, we consider signals generated from the string model in (3). It may here be noted that the MMLE is given by the non-linear least squares (NLS)

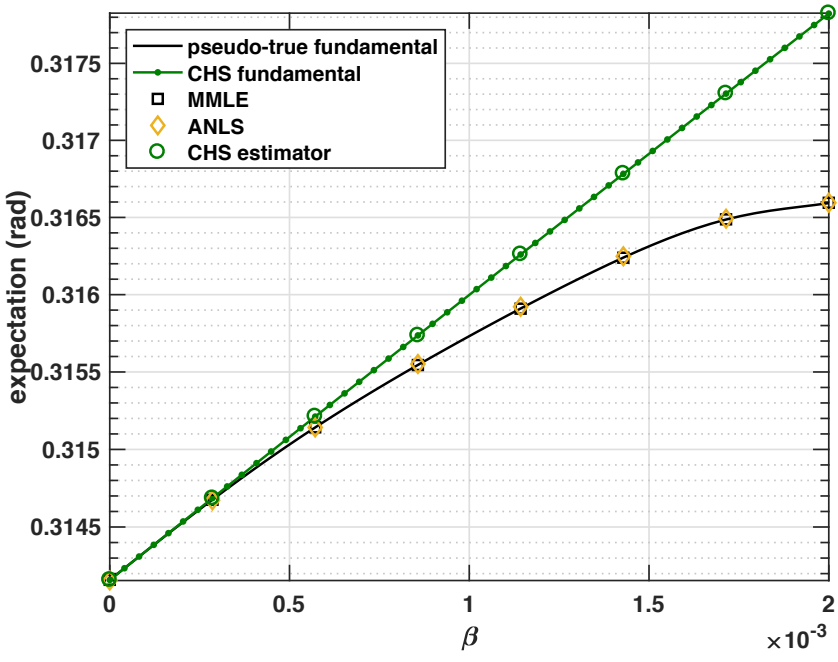

Figure 5: The pseudo-true fundamental frequency in Def. 2 as well as the CHS fundamental in Def. 4 for signals generated from the string model in (3) for varying β. Also plotted are the estimated expectations of the MMLE, ANLS, and CHS estimators.

estimator from [27]. In particular, let the signal consist of $K = 5$ components with amplitudes $\tilde{r}_k = e^{-\rho(k-K/2)^2}$ with $\rho = 0.2$, and let $\omega_0 = \pi/10$. The initial phases $\tilde{\varphi}_k$ are chosen uniformly random on $[-\pi, \pi)$. Furthermore, let the signal be observed at $N = 500$ time instances, and let the SNR, defined as $\text{SNR} = 10\log_{10}\sum_k \tilde{r}_k^2/\tilde{\sigma}^2$, be 10 dB. For this setting, Figure 5 presents the pseudo-true pitch i.e., the ω_0 defined in Def. 2, when varying the string stiffness parameter β on $[0, 2 \times 10^{-3}]$. Also presented is the CHS ω_0 in Def. 4. As may be noted, both definitions correspond to small perturbations of the pitch, with the CHS definition displaying a more linear behavior for a larger range of β. Figure 5 also displays the estimated expected values for the MMLE and CHS estimators, obtained from 2000 Monte Carlo simulations for each value of β. The CHS pitches are computed based on unstructured ML estimates of the component

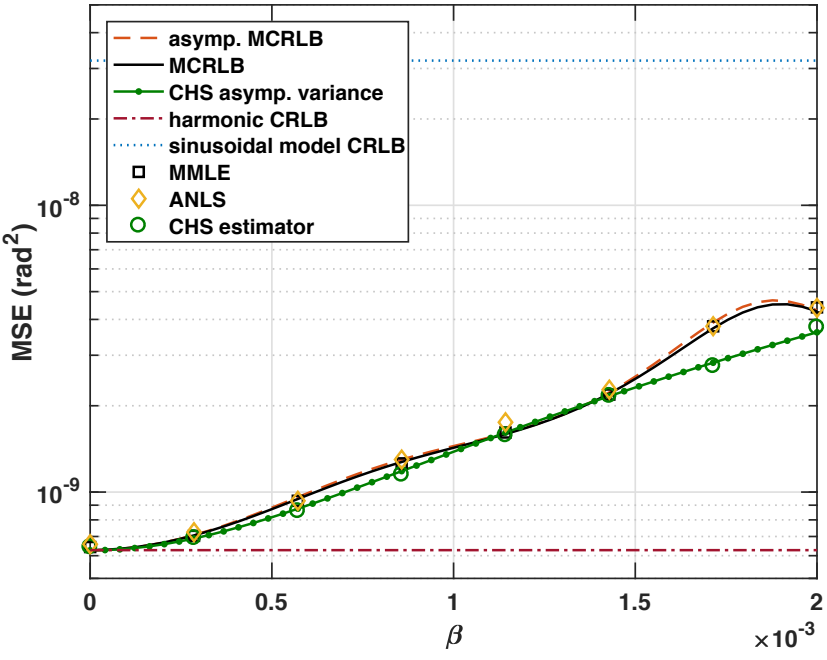

Figure 6: The MCRLB, the asymptotic MCRLB, and the asymptotic CHS estimator variance for estimating the pseudo-true ω_0 in Def. 2 and the CHS pitch in Def. 4, respectively, when the signal is generated from the string model in (3) for varying β. Also plotted are the corresponding MSE achieved by the MMLE, ANLS, and CHS estimators.

amplitudes and frequencies in accordance with Prop. 5. As can be seen, the sample averages correspond well to their theoretical values. Also presented is the average estimate obtained using the approximate NLS (ANLS) estimator [27], which is an asymptotic approximation of the MLE for an harmonic model (as $N \to \infty$), corresponding to harmonic summation from a periodogram estimate. It may here be noted that the expectation of the ANLS estimator coincides with the pseudo-true ω_0.

Furthermore, Figure 6 presents the exact and asymptotic MCRLB, as well the asymptotic variance of the CHS estimator, together with the MSE of the MMLE, ANLS, and CHS estimators. The MSE is here computed using the pseudo-true ω_0 as reference for the MMLE and ANLS, as it corresponds to their

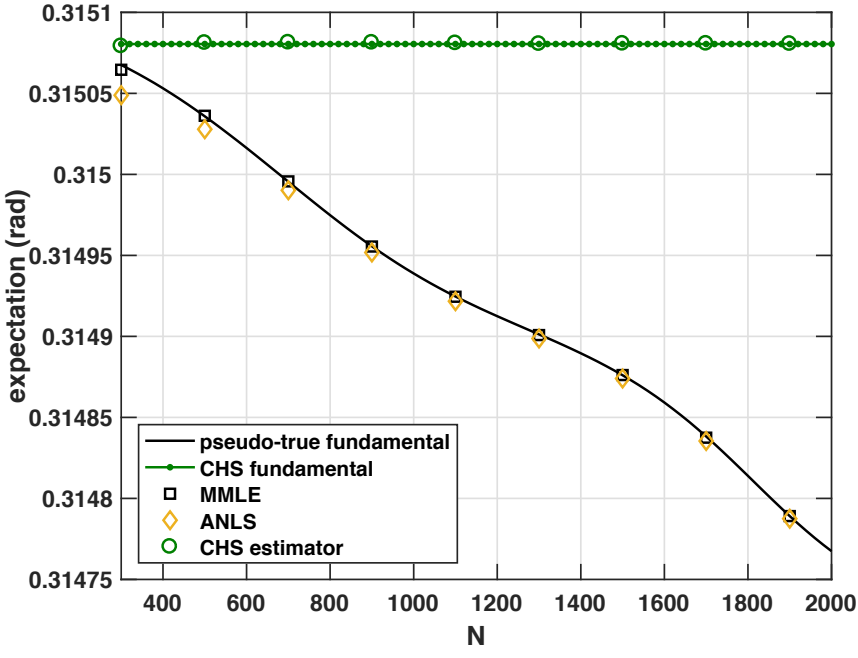

Figure 7: The pseudo-true fundamental frequency in Def. 2 as well as the CHS fundamental in Def. 4 for signals generated from the string model in (3) for varying N. Also plotted are the estimated expectations of the MMLE, ANLS, and CHS estimators.

expected and asymptotical expected values, respectively, whereas the reference for the CHS estimator is the CHS ω_0. As may be noted, the bounds provide accurate predictions of the behaviors of the three estimators, with the asymptotic MCRLB coinciding fairly well with its exact counterpart. As reference, the CRLB for the corresponding perfectly harmonic model, as well as for the lowest-frequency component in an unstructured sinusoidal model with exactly the same spectral content, are also provided. Figures 7 and 8 display corresponding quantities, i.e., theoretical and estimated expected values and theoretical and estimated variances, respectively, when fixing $\beta = 5 \times 10^{-4}$ and varying the number of samples, N, between 300 and 2000. As can be seen in Figures 7, the CHS ω_0 does not depend on the sample length, as expected. In contrast, the pseudo-true ω_0 depends on N. As $\tilde{r}_2 = \tilde{r}_3$ dominate the amplitudes of the

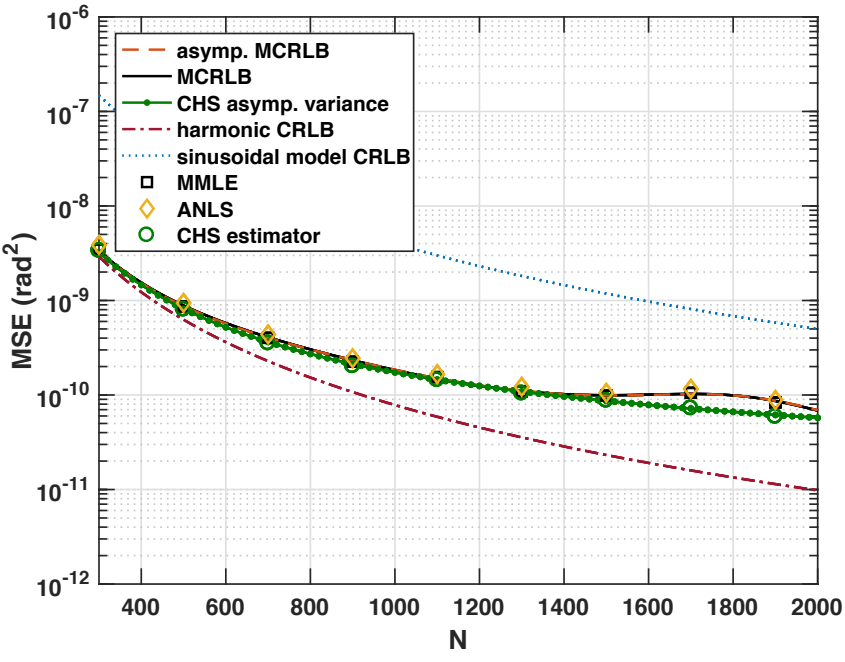

Figure 8: The MCRLB, the asymptotic MCRLB, and the asymptotic CHS estimator variance for estimating the pseudo-true ω_0 in Def. 2 and the CHS pitch in Def. 4, respectively, when the signal is generated from the string model in (3) for varying N. Also plotted are the corresponding MSE achieved by the MMLE, ANLS, and CHS estimators.

other components, we expect the ω_0 from Def. 2 to fluctuate close to the interval $[\omega_0\sqrt{1+\beta 2^2}, \omega_0\sqrt{1+\beta 3^2}] \approx [0.3145, 0.3149]$ for large but moderate values of N (recall that Def. 2 becomes ambiguous as $N \to \infty$). Furthermore, as may be seen from Figure 8, the MCRLB is far from being linear in the log of N, in contrast to the CHS asymptotic variance. It may here be noted that the slopes for the perfectly harmonic CRLB and the CHS asymptotic variance are different as the second term of (17) dominates for large N. For the same setting, Figure 9 presents the covariance bounds and obtained estimator MSEs for fixed $N = 500$, $\beta = 5 \times 10^{-4}$, and the SNR varying between -10 dB and 30 dB. As can be seen, the MCRLB and the CHS asymptotic variance are attained by the MMLE and the CHS estimator, respectively, for high enough SNR. In contrast, for the

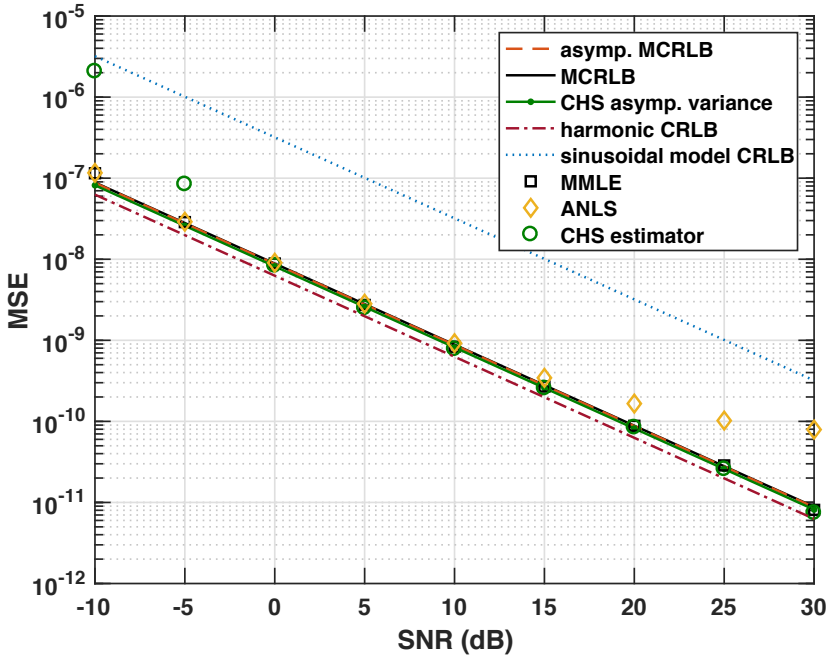

Figure 9: The MCRLB, the asymptotic MCRLB, and the asymptotic CHS estimator variance for estimating the pseudo-true ω_0 in Def. 2 and the CHS pitch in Def. 4, respectively, when the signal is generated from the string model in (3) for varying SNR. Also plotted are the corresponding MSE achieved by the MMLE, ANLS, and CHS estimators.

highest SNR values, the MCRLB is not attained by the ANLS estimator; as the ANLS is only an asymptotic approximation of the MMLE, the approximation error eventually becomes comparable to the noise variance, as the SNR tends to infinity.

6.2 Stochastic waveform

We proceed by extending the simulation study to the stochastic model in (18). Specifically, we let all signal parameters remain the same, with the difference being that each component frequency is perturbed by a random Gaussian number. Fixing $N = 500$ and SNR $= 10$ dB, we compute the HCRLB for ω_0, as well as for the (random) frequency of the first sinusoidal component, i.e., $\tilde{\omega}_1 = \omega_0 + \Delta_1$,

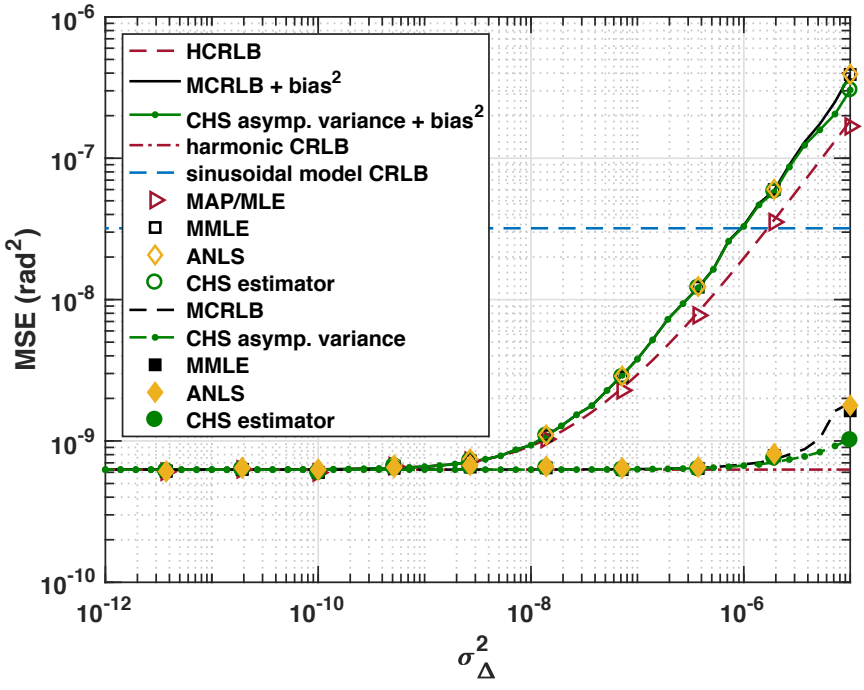

Figure 10: The MSE obtained when estimating ω_0 from the model in (18) for varying values of the inharmonicity parameter variance σ^2_Δ. Also given are MSEs for the estimates of the fundamental frequencies in Defs. 2 and 4.

for varying values of the inharmonicity variance σ^2_Δ. As comparison, we compute the theoretical MSE for the estimators implied by Defs. 2 and 4. This is performed by sampling $\Delta_k \in \mathcal{N}\left(0, \sigma^2_\Delta\right)$, computing the MSE for that particular set of $\{\Delta_k\}$, and averaging over realizations. Specifically, the MSE for a given $\{\Delta_k\}$ is computed by adding the squared bias, with references ω_0 and $\omega_0 + \Delta_1$, to the covariance bound, i.e., the MCRLB and the CHS asymptotic variance. For each value of σ^2_Δ, the corresponding MSEs are computed by averaging over 1000 Monte Carlo simulations. Furthermore, for each such simulation, noise is added to the signal waveform according to (5), and the signal parameters are estimated using MAP/MLE, MMLE, ANLS, and the CHS estimators. The results are displayed in Figures 10 and 11, with Figure 10 showing the MSE for the estimate of ω_0 and Figure 11 the MSE for $\tilde{\omega}_1 = \omega_0 + \Delta_1$. As may be seen from Figure 10, the

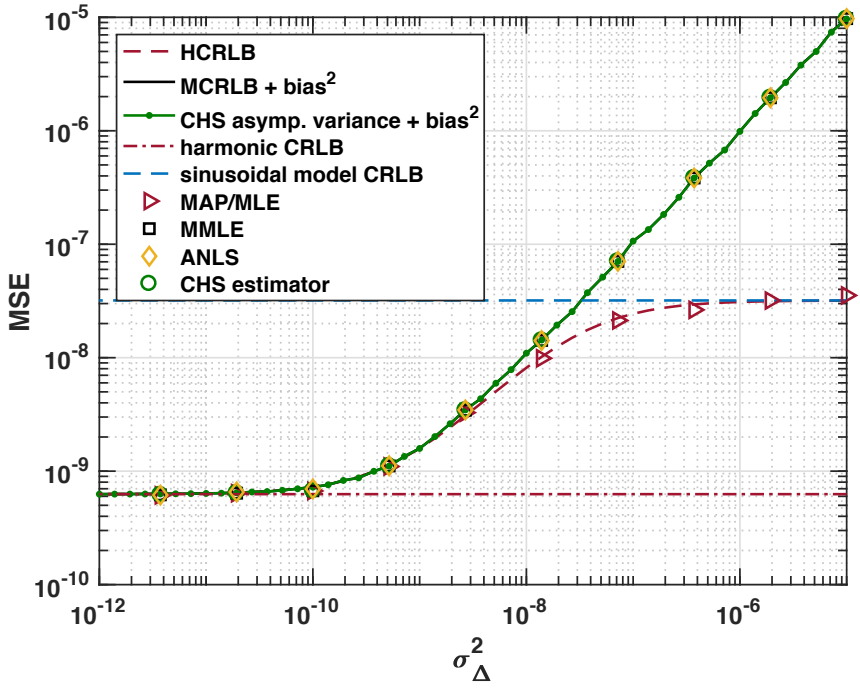

Figure 11: The MSE obtained when estimating the frequency $\tilde{\omega}_1 = \omega_0 + \Delta_1$ of the first sinusoidal component from the model in (18) for varying values of the inharmonicity parameter variance σ^2_Δ.

MAP/MLE estimator is able to attain the HCRLB, which is strictly smaller than the bounds for the estimators derived for the deterministic models. For reference, Figure 10 provides also the MCRLB, the theoretical CHS asymptotic variance, as well as the obtained MSEs for estimates of the ω_0 of Defs. 2 and 4. From this, it may be concluded that the variances of the estimators are dominated by their squared bias when estimating ω_0 from the model in (18). Furthermore, it may be noted from Figure 11 that when considering the estimate of the frequency of the first sinusoidal component, the HCRLB, as well as the MSE of the MAP/MLE, tends to the CRLB corresponding to an unstructured sinusoidal model as σ^2_Δ grows, which is due to negative correlation between the estimates of ω_0 and Δ_1. In contrast, for large enough inharmonicity, the MSEs of the MMLE, ANLS, and CHS estimators exceed the CRLB of the sinusoidal model.

7 Discussion

The three different definitions of the pitch for inharmonic signals all have their merits, and one cannot in a strictly objective sense say that one is better than the others. Instead, their relative usefulness depend on the considered application, as well as on ones aim with approximating close-to-harmonic signals with perfectly harmonic counterparts.

As noted, by defining pitch using an approximation in ℓ_2, as in Def. 2, one models the scenario of erroneously assuming that the observed waveform is periodic, and provides means for analyzing the average behavior of the misspecified MLE, as well as providing a bound on the estimation performance. Thus, Def. 2 is well suited for being used as a benchmark when analyzing the behavior of pitch estimators when applied to inharmonic signals; the definition makes it possible to compute the bias and MSE, at least empirically, for any such estimator. However, the behavior of Def. 2 is, as demonstrated in the numerical examples, non-linear with respect to both the inharmonicity and sample length, and is in addition ambiguous for very long signals.

In this respect, the OMT-based approximation of Def. 4 has the appealing quality of not depending on the actual signal, but only on its spectral properties. As shown both in the theoretical and numerical results, the behavior of Def. 4 is locally linear with respect to the signal inharmonicity, and corresponds well to the intuitive idea that inharmonicity corresponds to perturbations in frequency of spectral peaks. In addition to this, when used for estimation, Def. 4 displays linear behavior in terms of estimator variance, and coincides with the MLE for the perfectly harmonic case. With this in mind, Def. 4 is more satisfactory than Def. 2 when considered a tool for understanding inharmonic signals. Furthermore, its connection to the EXIP framework, i.e., the idea of fitting given parameter estimates to a certain structure, makes it relevant as a benchmark also in practical estimation scenarios.

In contrast to Defs. 2 and 4, which constitute approximations, Def. 5 views the inharmonic signal as a waveform resulting from perturbing the frequencies of the sinusoidal components by zero-mean random variables, causing them to deviate from perfect integer multiples of the nominal pitch. In this respect, Def. 5 offers an explanation of why the observed signal waveform is not periodic. However, as any given observation of the signal has been generated by this random procedure, one by definition cannot compute the ω_0 in Def. 5 from such an observation. However, adopting this view of inharmonic signals allows for com-

puting a performance bound constituting a smooth interpolant between perfectly harmonic models and completely unstructured sinusoidal models, as well as for finding an easily implemented estimator attaining the bound. In practical terms, Def. 5 offer tools for processing inharmonic signals for which one suspects that there may be no deterministic description of the inharmonic deviations. For example, such a class of signals could be the voiced part of human speech, for which one may observe frequency-dependent inharmonicity with no particular structure (this may very well change over time for a given person due to, e.g., infections). As the inharmonicity pattern thus could change from day to day for any given fundamental frequency, Def. 5 likely provides a pragmatic view.

8 Appendix

8.1 Information matrix for the HCRLB

Utilizing

$$\mathbb{E}_{\Delta_k}\left(\cos\left((\omega+\Delta_k)t+\tilde{\varphi}_k\right)\right) = e^{-\frac{1}{2}\sigma_\Delta^2 t^2}\cos\left(\omega t+\tilde{\varphi}_k\right)$$
$$\mathbb{E}_{\Delta_k}\left(\sin\left((\omega+\Delta_k)t+\tilde{\varphi}_k\right)\right) = e^{-\frac{1}{2}\sigma_\Delta^2 t^2}\sin\left(\omega t+\tilde{\varphi}_k\right),$$

one may write $\mathbb{E}_\Delta\left(F(\breve{\vartheta})\right) = \sum_{t=0}^{N-1}\Lambda^{(t)}$, where the matrices

$$\Lambda^{(t)} \triangleq \mathbb{E}_\Delta\left(\nabla_{\breve{\vartheta}}\mu_t^{\Re}(\breve{\vartheta})\nabla_{\breve{\vartheta}}\mu_t^{\Re}(\breve{\vartheta})^T + \nabla_{\breve{\vartheta}}\mu_t^{\Im}(\breve{\vartheta})\nabla_{\breve{\vartheta}}\mu_t^{\Im}(\breve{\vartheta})^T\right),$$

for $t = 0,\ldots,N-1$, may be expressed as

$$\Lambda^{(t)}_{\omega_0,\omega_0} = t^2\sum_k k^2\tilde{r}_k^2 + t^2 e^{-\sigma_\Delta^2 t^2}\sum_{(k,\ell):k\neq\ell} k\ell\tilde{r}_k\tilde{r}_\ell\cos\left(\omega(k-\ell)t+\tilde{\varphi}_k-\tilde{\varphi}_\ell\right)$$
$$\Lambda^{(t)}_{\omega_0,r_k} = te^{-\sigma_\Delta^2 t^2}\sum_{\ell:\ell\neq k}\ell\tilde{r}_\ell\sin\left(\omega(k-\ell)t+\tilde{\varphi}_k-\tilde{\varphi}_\ell\right)$$
$$\Lambda^{(t)}_{\omega_0,\varphi_k} = tk\tilde{r}_k^2 + te^{-\sigma_\Delta^2 t^2}\sum_{\ell:\ell\neq k}\ell\tilde{r}_k\tilde{r}_\ell\cos\left(\omega(k-\ell)t+\tilde{\varphi}_k-\tilde{\varphi}_\ell\right)$$
$$\Lambda^{(t)}_{r_k,r_\ell} = \begin{cases}1, & \ell = k\\ e^{-\sigma_\Delta^2 t^2}\cos\left(\omega(k-\ell)t+\tilde{\varphi}_k-\tilde{\varphi}_\ell\right), & \ell\neq k\end{cases}$$
$$\Lambda^{(t)}_{r_k,\varphi_\ell} = \begin{cases}0, & \ell = k\\ e^{-\sigma_\Delta^2 t^2}\sin\left(\omega(k-\ell)t+\tilde{\varphi}_k-\tilde{\varphi}_\ell\right), & \ell\neq k\end{cases}$$

$$\Lambda^{(t)}_{\varphi_k,\varphi_\ell} = \begin{cases} \tilde{r}_k^2, & \ell = k \\ e^{-\sigma_\Delta^2 t^2}\tilde{r}_k\tilde{r}_\ell \cos\big(\omega(k-\ell)t + \tilde{\varphi}_k - \tilde{\varphi}_\ell\big), & \ell \neq k \end{cases}$$

and $\Lambda^{(t)}_{\eta,\Delta_k} = t\Lambda^{(t)}_{\eta,\varphi_k}$, for $\eta \in \{\omega_0, r_k, \varphi_k\}$, and $\Lambda^{(t)}_{\Delta_k,\Delta_\ell} = t^2\Lambda^{(t)}_{\varphi_k,\varphi_\ell}$.

8.2 Proof of Theorem 1

Proof. Let p denote the pdf of the assumed model. Then, it can be shown that, for estimators $\hat{\vartheta}$ satisfying ML conditions, it holds that (see [13])

$$\mathbb{E}_{\mathbf{y}}\left((\hat{\vartheta} - \vartheta_0)(\hat{\vartheta} - \vartheta_0)^T\right) \succeq A(\vartheta_0)^{-1}F(\vartheta_0)A(\vartheta_0)^{-1},$$

with $\mathbb{E}_{\mathbf{y}}$ denoting expectation with respect to the pdf of the measured signal, and where

$$A(\vartheta_0) = \mathbb{E}_{\mathbf{y}}\left(\nabla^2_\vartheta p(\mathbf{y};\vartheta_0)\right),$$
$$F(\vartheta_0) = \mathbb{E}_{\mathbf{y}}\left(\nabla_\vartheta p(\mathbf{y};\vartheta_0)\nabla_\vartheta p(\mathbf{y};\vartheta_0)^T\right).$$

In the herein considered case, the model pdf is given by

$$p(\mathbf{y};\vartheta) = \frac{1}{(\pi\sigma^2)^N}\exp\left(-\frac{1}{\sigma^2}\sum_{t=0}^{N-1}|y_t - \mu_t(\vartheta)|^2\right)$$
$$= \frac{1}{(\pi\sigma^2)^N}\exp\left(-\frac{1}{\sigma^2}\sum_{t=0}^{N-1}(y_t^{\mathfrak{I}} - \mu_t^{\mathfrak{I}}(\vartheta))^2 + (y_t^{\mathfrak{R}} - \mu_t^{\mathfrak{R}}(\vartheta))^2\right).$$

As the additive noise of the actual measured signal is circularly symmetric white Gaussian noise with variance $\tilde{\sigma}^2$, it is readily verified that

$$\mathbb{E}_{\mathbf{y}}\left(\frac{\partial p(\mathbf{y};\vartheta)}{\partial\vartheta_k}\frac{\partial p(\mathbf{y};\vartheta)}{\partial\vartheta_\ell}\right) = B_{k,\ell}(\vartheta) + \tilde{B}_{k,\ell}(\vartheta)$$

where k and ℓ denotes indices of the elements of ϑ and

$$B_{k,\ell}(\vartheta) = \frac{2\tilde{\sigma}^2}{(\sigma^2)^2}\sum_{t=0}^{N-1}\frac{\partial\mu_t^{\mathfrak{R}}(\vartheta)}{\partial\vartheta_k}\frac{\partial\mu_t^{\mathfrak{R}}(\vartheta)}{\partial\vartheta_\ell} + \frac{\partial\mu_t^{\mathfrak{I}}(\vartheta)}{\partial\vartheta_k}\frac{\partial\mu_t^{\mathfrak{I}}(\vartheta)}{\partial\vartheta_\ell}$$

and

$$\tilde{B}_{k,\ell}(\vartheta) = \frac{4}{(\sigma^2)^2}\left(\sum_{t=0}^{N-1}\frac{\partial\mu_t^{\mathfrak{R}}(\vartheta)}{\partial\vartheta_k}\xi_t^{\mathfrak{R}}(\vartheta) + \frac{\partial\mu_t^{\mathfrak{I}}(\vartheta)}{\partial\vartheta_k}\xi_t^{\mathfrak{I}}(\vartheta)\right) \cdot \left(\sum_{s=0}^{N-1}\frac{\partial\mu_s^{\mathfrak{R}}(\vartheta)}{\partial\vartheta_\ell}\xi_s^{\mathfrak{R}}(\vartheta) + \frac{\partial\mu_s^{\mathfrak{I}}(\vartheta)}{\partial\vartheta_\ell}\xi_s^{\mathfrak{I}}(\vartheta)\right),$$

with ξ given in Proposition 1. However, by definition, the pseudo-true parameter ϑ_0 solves the least-squares problem in (6), and in particular, ϑ_0 satisfies the first-order optimality conditions, i.e.,

$$\sum_{t=0}^{N-1}\frac{\partial\mu_t^{\mathfrak{R}}(\vartheta_0)}{\partial\vartheta_k}\xi_t^{\mathfrak{R}}(\vartheta_0) + \frac{\partial\mu_t^{\mathfrak{I}}(\vartheta_0)}{\partial\vartheta_k}\xi_t^{\mathfrak{I}}(\vartheta_0) = 0$$

for all parameter indices k. Thus, $\tilde{B}_{k,\ell}(\vartheta_0) = 0$ for all k and ℓ, yielding the expression for $F(\vartheta_0)$. Furthermore, it is straightforward to verify that

$$-\mathbb{E}_{\mathbf{y}}\left(\frac{\partial^2 p(\mathbf{y};\vartheta)}{\partial\vartheta_k\partial\vartheta_\ell}\right) = \frac{2}{\sigma^2}\sum_{t=0}^{N-1}\frac{\partial\mu_t^{\mathfrak{R}}(\vartheta)}{\partial\vartheta_k}\frac{\partial\mu_t^{\mathfrak{R}}(\vartheta)}{\partial\vartheta_\ell} + \frac{\partial\mu_t^{\mathfrak{I}}(\vartheta)}{\partial\vartheta_k}\frac{\partial\mu_t^{\mathfrak{I}}(\vartheta)}{\partial\vartheta_\ell} - \frac{2}{\sigma^2}\sum_{t=0}^{N-1}\frac{\partial^2\mu_t^{\mathfrak{R}}(\vartheta)}{\partial\vartheta_k\partial\vartheta_\ell}\xi_t^{\mathfrak{R}}(\vartheta) + \frac{\partial^2\mu_t^{\mathfrak{I}}(\vartheta)}{\partial\vartheta_k\partial\vartheta_k}\xi_t^{\mathfrak{I}}(\vartheta),$$

for any ϑ, which yields the expression for $A(\vartheta_0)$. □

8.3 Proof of Proposition 2

Proof. By Lemma 1 below, $\frac{1}{N}F(\vartheta_0)$ and $\frac{1}{N}A(\vartheta_0)$ converge to arrowhead matrices. As shorthands, let $F = F(\vartheta_0)$ and $A = A(\vartheta_0)$. For large N, the structure of $\frac{1}{N}A$ is

$$\frac{1}{N}A = -\frac{1}{N\sigma^2}\begin{bmatrix}\eta & z^T\\ z & \operatorname{diag}(d)\end{bmatrix},$$

where $\operatorname{diag}(d)$ denotes the diagonal matrix with the vector d as its main diagonal, and where $\eta = \eta_\mu + \eta_\xi$ and $z = z_\mu + z_\xi$ with

$$\eta_\mu = 2\sum_{t=0}^{N-1}\left(\frac{\partial\mu_t^{\mathfrak{R}}}{\partial\omega}\right)^2 + \left(\frac{\partial\mu_t^{\mathfrak{I}}}{\partial\omega}\right)^2$$

$$\eta_\xi = 2\sum_{t=0}^{N-1} \xi_t^{\Re}\frac{\partial^2 \mu_t^{\Re}}{\partial \omega^2} + \xi_t^{\Im}\frac{\partial^2 \mu_t^{\Im}}{\partial \omega^2}$$

$$d = 2\sum_{t=0}^{N-1} \nabla_\alpha \mu_t^{\Re} \odot \nabla_\alpha \mu_t^{\Re} + \nabla_\alpha \mu_t^{\Im} \odot \nabla_\alpha \mu_t^{\Im}$$

$$z_\mu = 2\sum_{t=0}^{N-1} \nabla_\alpha \mu_t^{\Re}\frac{\partial \mu_t^{\Re}}{\partial \omega} + \nabla_\alpha \mu_t^{\Im}\frac{\partial \mu_t^{\Im}}{\partial \omega}$$

$$z_\xi = 2\sum_{t=0}^{N-1} \xi_t^{\Re}\frac{\partial}{\partial \omega}\nabla_\alpha \mu_t^{\Re} + \xi_t^{\Im}\frac{\partial}{\partial \omega}\nabla_\alpha \mu_t^{\Im},$$

where $\odot$ denotes the Hadamard product, and

$$\alpha = \begin{bmatrix} \varphi_1 & \dots & \varphi_K & r_1 & \dots & r_K \end{bmatrix}^T,$$

with all derivatives being evaluated at $\vartheta = \vartheta_0$. Then, by the Sherman-Morrison-Woodbury formula [28], A^{-1} may be well approximated as

$$A^{-1} = -\sigma^2 \begin{bmatrix} 0 & 0 \\ 0 & \operatorname{diag}(d)^{-1} \end{bmatrix} - \frac{\sigma^2}{\rho} uu^T$$

for large N, where $u = u_\mu + u_\xi$, with $\rho = \eta - z^T(z./d)$ and

$$u_\mu = \begin{bmatrix} -1 & (z_\mu./d)^T \end{bmatrix}^T,\ u_\xi = \begin{bmatrix} 0 & (z_\xi./d)^T \end{bmatrix}^T,$$

where ./ denotes elementwise division. Similarly, F may be approximated as

$$F = \frac{\tilde{\sigma}^2}{(\sigma^2)^2}\begin{bmatrix} \eta_\mu & z_\mu^T \\ z_\mu & \operatorname{diag}(d) \end{bmatrix}$$

for large N. Then, the MCRLB corresponding to ω is given by the first diagonal element of $A^{-1}FA^{-1}$, which is given by

$$\frac{(\sigma^2)^2}{\rho^2} u^T F u = (\sigma^2)^2 \frac{1}{\rho^2}\left(u_\mu^T F u_\mu + u_\xi^T F u_\xi\right) = \sigma^2 \frac{C+E}{(C-E+Z+D)^2},$$

where it used that $u_\mu^T F u_\xi = 0$, and

$$C = \eta_\mu - z_\mu^T(z_\mu./d)\ ,\ E = z_\xi^T(z_\xi./d)\ ,\ D = -2z_\mu^T(z_\xi./d)\ ,\ Z = \eta_\xi.$$

Assuming that the pseudo-true fundamental frequency is not too close to zero, the correlation between signal components corresponding to different harmonic orders tends to zero as $N \to \infty$. The asymptotic expressions for C, E, D, and Z stated in the proposition follow directly. □

Lemma 1. *As $N \to \infty$, $\frac{1}{N}F(\vartheta_0)$ and $\frac{1}{N}A(\vartheta_0)$ converge to arrowhead matrices.*

Proof. Firstly, it may be noted that as ϑ_0 solves the least squares criterion in (6), it directly follows from the optimality criterion that

$$\sum_{t=0}^{N-1} \xi_t^{\Re} \nabla_{\vartheta[l]} \mu_t^{\Re} + \sum_{t=0}^{N-1} \xi_t^{\Im} \nabla_{\vartheta[l]} \mu_t^{\Im} = 0.$$

Then, as any second derivative of $\mu_t^{\Re}$ and $\mu_t^{\Im}$ not involving differentiation with respect to ω is equal to a common constant real scaling of elements of $\nabla_{\vartheta[l]}\mu_t^{\Re}$ and $\nabla_{\vartheta[l]}\mu_t^{\Im}$, respectively, only the first column and first row of $\tilde{F}(\vartheta_0)$ are non-zero. It is straightforward to show that elements of $F(\vartheta_0)$ not related to partial derivatives of ω converge linearly to zero when scaled by $1/N$, whereas the diagonal is bounded from below by positive values. Thus, $\frac{1}{N}F(\vartheta_0)$, and thereby $\frac{1}{N}A(\vartheta_0)$, converges to an arrowhead matrix as $N \to \infty$. □

8.4 Proof of Proposition 4

Proof. Clearly, if there are at least two consecutive sinusoids with non-zero amplitude, d in Definition 3 satisfies $d \in \left[\tilde{\omega}_0 - 2\left\|\Delta\right\|_\infty, \tilde{\omega}_0 + 2\left\|\Delta\right\|_\infty\right]$. Then,

$$(K+1)d \geq \tilde{\omega}_K \ , \ (K-1)d < \tilde{\omega}_K,$$

if $\left\|\Delta\right\|_\infty \leq \tilde{\omega}_0/(2K+3)$. Thus for $\left\|\Delta\right\|_\infty < \tilde{\omega}_0/(2K+3)$, we have that $L \in \{K, K+1\}$. Furthermore, for any $\hat{\omega} \in \left[\omega - \left\|\Delta\right\|_\infty, \omega + \left\|\Delta\right\|_\infty\right]$,

$$|k\hat{\omega} - \tilde{\omega}_k| \leq (k+1)\left\|\Delta\right\|_\infty = \tilde{\omega}_0 \frac{k+1}{2K+3},$$

whereas

$$|(k \pm 1)\hat{\omega} - \tilde{\omega}_k| \geq \tilde{\omega}_0 \frac{2K+1-k}{2K+3}.$$

Thus, $\arg\min_{\ell} (\ell\hat{\omega} - \tilde{\omega}_k)^2 = k$, for $k = 1, \ldots, K$, implying

$$q_L(\hat{\omega}) = 2\pi \sum_{k=1}^{K} \tilde{r}_k^2 (k\hat{\omega} - \tilde{\omega}_k)^2 = 2\pi \sum_{k=1}^{K} \tilde{r}_k^2 (k(\hat{\omega} - \tilde{\omega}_0) - \Delta_k)^2$$

for any $\hat{\omega} \in \left[\tilde{\omega}_0 - \|\Delta\|_\infty, \tilde{\omega}_0 + \|\Delta\|_\infty\right]$. This quadratic function has the unique stationary point

$$\omega_0 = \frac{\sum_{k=1}^{K} \tilde{r}_k^2 k\tilde{\omega}_k}{\sum_{k=1}^{K} \tilde{r}_k^2 k^2} = \tilde{\omega}_0 + \frac{\sum_{k=1}^{K} \tilde{r}_k^2 k\Delta_k}{\sum_{k=1}^{K} \tilde{r}_k^2 k^2},$$

where it may be noted that

$$\left| \frac{\sum_{k=1}^{K} \tilde{r}_k^2 k\Delta_k}{\sum_{k=1}^{K} \tilde{r}_k^2 k^2} \right| \leq \frac{\sum_{k=1}^{K} \tilde{r}_k^2 k}{\sum_{k=1}^{K} \tilde{r}_k^2 k^2} \|\Delta\|_\infty \leq \|\Delta\|_\infty .$$

□

8.5 Proof of Proposition 5

Proof. We here assume that the inharmonic perturbations are small so that, asymptotically, i.e., when $N \to \infty$ or the SNR tending to infinity, the assumptions of Proposition 4 hold in the sense that there exist ω such that $|k\omega - \hat{\tilde{\omega}}_k| \leq \omega/(2K+3)$, for $k = 1, \ldots, K$, almost surely. Noting that the covariance matrix of the vector $\hat{\vartheta} = \begin{bmatrix} \hat{\tilde{r}}_1 & \ldots & \hat{\tilde{r}}_K & \hat{\tilde{\omega}}_1 & \ldots & \hat{\tilde{\omega}}_K \end{bmatrix}^T$ is asymptotically given by (see, e.g., [29])

$$\mathrm{Cov}(\hat{\vartheta}) = \frac{\tilde{\sigma}^2}{2N} \begin{bmatrix} I & 0 \\ 0 & C_2 \end{bmatrix}, \; C_2 = \frac{12}{N^2 - 1} \mathrm{diag}\left(\begin{bmatrix} 1/\tilde{r}_1^2 & \ldots & 1/\tilde{r}_K^2 \end{bmatrix}\right)$$

and that the estimate of the CHS fundamental frequency is given $\hat{\omega}_0 = f(\hat{\vartheta})$, where

$$f(\vartheta) = \frac{\sum_{k=1}^{K} \tilde{r}_k^2 k\tilde{\omega}_k}{\sum_{k=1}^{K} \tilde{r}_k^2 k^2},$$

the expression in (17) is the obtained from a first order Taylor expansion of f at ϑ. As can be seen from $\mathrm{Cov}(\hat{\vartheta})$, the estimates of $\tilde{\omega}_k$ are asymptotically uncorrelated

with estimates of the amplitudes, from which it directly follows that $\hat{\omega}_0$ is an asymptotically unbiased estimate of ω_0. As $\hat{\vartheta}$ is the MLE of ϑ, consistency follows, with the asymptotic variance being given by

$$\mathbb{E}\left((\omega_0 - \hat{\omega}_0)^2\right) = \nabla_\vartheta f(\vartheta)^T \mathrm{Cov}(\hat{\vartheta}) \nabla_\vartheta f(\vartheta).$$

After some simplification, the expression in (17) follows. □

8.6 Proof of Proposition 7

Proof. We have

$$\nabla_{\breve{\vartheta}} \log p(\mathbf{y}, \Delta; \vartheta) = \nabla_{\breve{\vartheta}} \log p(\mathbf{y} \mid \Delta; \vartheta) + \nabla_{\breve{\vartheta}} \log p(\Delta),$$

where $\nabla_{\breve{\vartheta}} \log p(\Delta) = \mathrm{diag}\begin{pmatrix} 0^T & -\Delta^T/\sigma_\Delta^2 \end{pmatrix}$, where 0 is a zero vector of length $2K+1$. Further, as

$$\mathbb{E}_{\mathbf{y}|\Delta}\left(\nabla_{\breve{\vartheta}} \log p(\mathbf{y} \mid \Delta; \vartheta)\right) = 0,$$

and Δ is independent of the measurement noise, it follows that

$$\mathbb{E}_{\mathbf{y}|\Delta}\left(\nabla_{\breve{\vartheta}} \log p \; \nabla_{\breve{\vartheta}} \log p^T\right) = F(\breve{\vartheta}) + \begin{bmatrix} 0 & 0 \\ 0 & \frac{1}{(\sigma_\Delta^2)^2}\Delta\Delta^T \end{bmatrix},$$

where we use the shorthand $p = p(\mathbf{y}, \Delta; \vartheta)$, and

$$\begin{aligned} F(\breve{\vartheta}) &= \mathbb{E}_{\mathbf{y}|\Delta}\left(\nabla_{\breve{\vartheta}} \log p(\mathbf{y} \mid \Delta; \vartheta) \nabla_{\breve{\vartheta}} \log p(\mathbf{y} \mid \Delta; \vartheta)^T\right) \\ &= \frac{2}{\tilde{\sigma}^2} \sum_{t=0}^{N-1} \nabla_{\breve{\vartheta}} \mu_t^{\Re}(\breve{\vartheta}) \nabla_{\breve{\vartheta}} \mu_t^{\Re}(\breve{\vartheta})^T + \nabla_{\breve{\vartheta}} \mu_t^{\Im}(\breve{\vartheta}) \nabla_{\breve{\vartheta}} \mu_t^{\Im}(\breve{\vartheta})^T \end{aligned}$$

as the measurement noise is circularly symmetric white Gaussian. As $\mathbb{E}_\Delta\left(\Delta\Delta^T\right) = \sigma_\Delta^2 I$, the expression for $\breve{F}$ follows directly. It may also be readily verified that

$$\mathbb{E}_{\mathbf{y}|\Delta}\left(\nabla_{\breve{\vartheta}} \log p(\mathbf{y} \mid \Delta; \vartheta) \frac{\partial}{\partial \tilde{\sigma}^2} \log p(\mathbf{y} \mid \Delta; \vartheta)\right) = 0,$$

i.e., the HCRLB for $\breve{\vartheta}$ does not depend on whether $\tilde{\sigma}^2$ is known or not, implying that no partial derivatives with respect to $\tilde{\sigma}^2$ need to be considered in order to compute the HCRLB of $\breve{\vartheta}$. □

8.7 Proof of Proposition 8

Proof. Clearly, for any choice of $\breve{\vartheta}$, $\mathcal{L}$ is maximized with respect to $\tilde{\sigma}^2$ by

$$\tilde{\sigma}^2 = \frac{1}{N} \left\| \mathbf{y} - \mathbf{x}(\breve{\vartheta}) \right\|_2^2 .$$

Plugging this in, we get, excluding constant terms, the criterion

$$\mathcal{L}' = -N \log \|\mathbf{y} - \mathbf{x}(\vartheta, \Delta)\|_2^2 - \frac{1}{2\sigma_\Delta^2} \|\Delta\|_2^2 .$$

Then, letting $\boldsymbol{\omega} = \begin{bmatrix} \omega_0 + \Delta_1 & \dots & \omega_0 + \Delta_K \end{bmatrix}^T$, the criterion is, for any choice of (ω_0, Δ), minimized with respect to the complex amplitudes by

$$\mathbf{x}(\vartheta) = A(\boldsymbol{\omega}) \left(A(\boldsymbol{\omega})^H A(\boldsymbol{\omega}) \right)^{-1} A(\boldsymbol{\omega})^H \mathbf{y}.$$

Also, for any set of frequencies $\omega_1, \dots, \omega_K$, one may maximize the second term in the criterion, i.e., the negative of the sum of squared inharmonicities, by choosing the nominal fundamental frequency ω_0 as

$$\omega_0 = \frac{\sum_{k=1}^K k\omega_k}{\sum_{k=1}^K k^2},$$

which does not affect the data fit term $\left\| \mathbf{y} - \mathbf{x}(\breve{\vartheta}) \right\|_2^2$. This results in the criterion $\psi_{\mathrm{ML/MAP}}$, which is to be maximized over a set of K unconstrained frequencies. □

References

[1] S. M. Nørholm, J. R. Jensen, and M. G. Christensen, "Instantaneous Fundamental Frequency Estimation With Optimal Segmentation for Non-stationary Voiced Speech," *IEEE/ACM Trans. Audio, Speech, Language Process.*, vol. 24, no. 12, pp. 2354–2367, Dec 2016.

[2] M. Müller, D. P. W. Ellis, A. Klapuri, and G. Richard, "Signal Processing for Music Analysis," *IEEE J. Sel. Topics Signal Process.*, vol. 5, no. 6, pp. 1088–1110, 2011.

[3] R. B. Randall, *Vibration-Based Condition Monitoring: Industrial, Aerospace and Automotive Applications*, John Wiley & Sons, Chichester, UK, 2011.

[4] M. A. Little, P. E. McSharry, E. J. Hunter, J. Spielman, and L. O. Ramig, "Suitability of Dysphonia Measurements for Telemonitoring of Parkinson's disease," *IEEE Trans. Biomed. Eng.*, vol. 56, no. 4, pp. 1015–102, April 2009.

[5] J. K. Nielsen, T. L. Jensen, J. R. Jensen, M. G. Christensen, and S. H. Jensen, "Fast fundamental frequency estimation: Making a ststatistical efficient estimator computationally efficient," *Elsevier Signal Processing*, vol. 135, pp. 188–197, Jan 2017.

[6] M. Christensen and A. Jakobsson, *Multi-Pitch Estimation*, Morgan & Claypool, San Rafael, Calif., 2009.

[7] N. H. Fletcher and T. D. Rossing, *The Physics of Musical Instruments*, Springer-Verlag, New York, NY, 1988.

[8] E. B. George and M. J. T. Smith, "Speech analysis/synthesis and modification using an analysis-by-synthesis/overlap-add sinusoidal model," *IEEE Trans. Speech Audio Process.*, vol. 5, no. 5, pp. 389–406, Sep 1997.

[9] J. X. Zhang, M. G. Christensen, S. H. Jensen, and M. Moonen, "A Robust and Computationally Efficient Subspace-Based Fundamental Frequency Estimator," *IEEE Trans. Audio, Speech, Language Process*, vol. 18, no. 3, pp. 487–497, March 2010.

[10] N. R. Butt, S. I. Adalbjörnsson, S. D. Somasundaram, and A. Jakobsson, "Robust Fundamental Frequency Estimation in the Presence of Inharmonicities," in *38th IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, Vancouver, May 26–31, 2013.

[11] S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory*, Prentice-Hall, Englewood Cliffs, N.J., 1993.

[12] S. Fortunati, F. Gini, M. S. Greco, and C. D. Richmond, "Performance bounds for parameter estimation under misspecified models: Fundamental findings and applications," *IEEE Signal Process. Mag.*, vol. 34, no. 6, pp. 142–157, Nov 2017.

[13] C. D. Richmond and L. L. Horowitz, "Parameter Bounds on Estimation Accuracy Under Model Misspecification," *IEEE Trans. Signal. Process*, vol. 63, no. 9, pp. 2263–2278, 2015.

[14] C. Villani, *Optimal transport: old and new*, Springer Science & Business Media, 2008.

[15] P. Stoica and T. Söderström, "On Reparametrization of Loss Functions Used in Estimation and the Invariance Principle," *Signal Processing*, vol. 17, pp. 383–387, August 1989.

[16] A. Swindlehurst and P. Stoica, "Maximum Likelihood Methods in Radar Array Signal Processing," *IEEE Proc.*, vol. 86, no. 2, pp. 421–441, February 1998.

[17] H. Li, P. Stoica, and J. Li, "Computationally Efficient Parameter Estimation for Harmonic Sinusoidal Signals," *Signal Processing*, vol. 80, pp. 1937–1944, 2000.

[18] S. L. Marple, "Computing the discrete-time "analytic" signal via FFT," *IEEE Trans. Signal Process.*, vol. 47, no. 9, pp. 2600–2603, September 1999.

[19] A. Klapuri, "Multiple fundamental frequency estimation based on harmonicity and spectral smoothness," *IEEE Trans. Speech Audio Process.*, vol. 11, no. 6, pp. 804–816, 2003.

[20] F. Elvander, T. Kronvall, S. I. Adalbjörnsson, and A. Jakobsson, "An Adaptive Penalty Multi-Pitch Estimator with Self-Regularization," *Elsevier Signal Processing*, vol. 127, pp. 56–70, October 2016.

[21] R. J. McAulay and T. F. Quatieri, "Pitch estimation and voicing detection based on a sinusoidal speech model," in *Proc. 1990 Int. Conf. on Acoustics, Speech, and Signal Processing*, Albuquerque, NM, USA, 1990, pp. 249–252.

[22] M. G. Christensen, A. Jakobsson, and S. H. Jensen, "Joint High-Resolution Fundamental Frequency and Order Estimation," *IEEE Trans. Audio, Speech, Language Process*, vol. 15, no. 5, pp. 1635–1644, July 2007.

[23] T. T. Georgiou, J. Karlsson, and M. S. Takyar, "Metrics for power spectra: an axiomatic approach," *IEEE Trans. Signal Process.*, vol. 57, no. 3, pp. 859–867, Mar. 2009.

[24] H. Messer, "The Hybrid Cramer-Rao Lower Bound - From Practice to Theory," in *Fourth IEEE Workshop on Sensor Array and Multichannel Processing, 2006.*, 12–14 July, 2006, pp. 304–307.

[25] Y. Naim and H. Messer, "Notes on the Tightness of the Hybrid Cramér-Rao Lower Bound," *IEEE Trans. Signal Process.*, vol. 57, no. 6, pp. 2074–2084, Jun. 2009.

[26] A. Yeredor, "The Joint MAP-ML Criterion and its Relation to ML and to Extended Least-Squares," *IEEE Trans. Signal Process.*, vol. 48, no. 12, pp. 3484–3492, Dec. 2000.

[27] M. G. Christensen, P. Stoica, A. Jakobsson, and S. H. Jensen, "Multi-pitch estimation," *Signal Processing*, vol. 88, no. 4, pp. 972–983, April 2008.

[28] G. H. Golub and C. F. Van Loan, *Matrix Computations*, The John Hopkins University Press, 3^{rd} edition, 1996.

[29] P. Stoica, A. Jakobsson, and J. Li, "Cisoid Parameter Estimation in the Colored Noise Case: Asymptotic Cramér-Rao Bound, Maximum Likelihood and Nonlinear Least-Squares," *IEEE Trans. Signal Process.*, vol. 45, pp. 2048–2059, August 1997.

D

Paper D

Mismatched Estimation of Polynomially Damped Signals

Filip Elvander, Johan Swärd, and Andreas Jakobsson

Centre for Mathematical Sciences, Lund University, Lund, Sweden

Abstract

In this work, we consider the problem of estimating the parameters of polynomially damped sinusoidal signals, commonly encountered in, for instance, spectroscopy. Generally, finding the parameter values of such signals constitutes a high-dimensional problem, often further complicated by not knowing the number of signal components or their specific signal structures. In order to alleviate the computational burden, we herein propose a mismatched estimation procedure using simplified, approximate signal models. Despite the approximation, we show that such a procedure is expected to yield predictable results, allowing for statistically and computationally efficient estimates of the signal parameters.

Key words: Mismatched estimation, computational efficiency, NMR spectroscopy, Lorentzian and Voigt line shapes

1 Introduction

Signals that may be well modeled as a superposition of exponentially decaying complex-valued sinusoids appear in a wide variety of fields, such as radar, geology, non-destructive testing, and spectroscopy (see, e.g., [1–3]). In this work, we are primarily interested in nuclear magnetic resonance (NMR) spectroscopy, where the signal parameters correspond to properties of the material under study, such as intra-molecular forces. Historically, the model most commonly considered is the so-called Lorentzian line shape [4], i.e., wherein the decay of the signal components are modeled as an exponential first-degree polynomial, although more detailed signal models are also common, such as the Voigt line shape, which uses a second-order polynomial decay [5].

The estimation of the parameters of the Lorentzian signal model has been approached in a variety of ways (see [1] for a more general review), e.g., by exploiting subspace decompositions [6, 7], linear system descriptions [8], as well as compressed sensing methods [9–11]. However, for some responses, the first-order polynomial is insufficient for accurately modeling the observed data, and one instead requires the use of Voigt line shapes to more accurately capture the structure of the signal. This model extension then implies that methods based on linear prediction no longer are applicable, demanding more advanced procedures for estimating the signal parameters.

A common approach for addressing this issue is to form estimates by minimizing a non-linear least squares criterion [12], although such an approach requires prior information about the model order, or that such information is estimated, e.g., by adding sparsity enhancing penalties to the cost function, iteratively solving for one signal component at the time [13], or performing model order estimation using some predefined criterion [14]. Performing these searches can be computationally cumbersome due to the large parameter space that the signal model entail. To find a remedy for this problem, we propose an estimation procedure that gradually increases the model complexity, such that one is restricting the parameter space over which the parameter search is performed. To alleviate this, we use the framework mismatched estimation (see, e.g., [15] for a recent overview). Specifically, we present a brief analysis on the expected behavior of approximate maximum likelihood estimators (MLEs) derived under simplified model assumptions, i.e., when using lower-order polynomials for describing the signal decay.

In the first part of the paper, we describe the signal model and state the sought optimization problem. After this, we present the expected behavior of

mismatched MLEs for simplifications of this model. We then propose a computationally efficient statistical test based on the spectral properties for discriminating between different polynomial decay models. In particular, the proposed test does not require estimates of the parameters corresponding to the more complex signal models, thereby avoiding computationally cumbersome searches for redundant signal parameters. Lastly, we present numerical examples illustrating the performance of the proposed procedure as compared to the Cramér-Rao lower bound (CRLB) [13].

2 Signal model

Consider a signal consisting of K polynomially damped sinusoids[1]

$$y_n = \sum_{k=1}^{K} r_k e^{i\varphi_k + i\omega_k t_n - \beta_k t_n - \gamma_k t_n^2} + \varepsilon_n, \tag{1}$$

for $n = 0, \ldots, N-1$, where r_k denotes the amplitude, φ_k the phase, ω_k the angular frequency, $\beta_k \geq 0$ the Lorentzian damping, $\gamma_k \geq 0$ the Voigt damping for kth signal component, and t_n the time at sample n. Furthermore, ε_n denotes an additive noise, herein assumed to be well modeled as a circularly symmetric white Gaussian noise[2]. In this work, we aim to formulate a procedure for estimating the parameters of (1) in a computationally and statistically efficient manner. It should be stressed that no knowledge of K, nor of the number of components of each signal class, i.e., whether certain parameters β_k or γ_k are strictly positive, is assumed. Thus, the number of signal components, as well as their class, has to be estimated.

As a preliminary, it may be noted that the MLE, assuming knowledge of K, may be formed as

$$\underset{\psi_1,\ldots,\psi_K}{\arg\min} \sum_{n=0}^{N-1} \left| y_n - \sum_{k=1}^{K} \mu(t_n; \psi_k, \alpha) \right|^2, \tag{2}$$

[1] We here restrict our attention to Lorentzian and Voigt line shapes, but note that the presented procedure may also be used for higher order polynomial decays. Such models are used in, for instance, non-destructive testing [2].

[2] It may be noted that this constitutes a valid assumption in many spectroscopy applications, with ε_n corresponding to thermal (Johnson) noise.

where $\psi_k = [r_k, \varphi_k, \omega_k, \beta_k, \gamma_k]$ is the parameter vector for component k, and

$$\mu(t_n; \psi_k, \alpha) = r_k \alpha(t_n; \psi_k) e^{i\varphi_k + i\omega_k t_n} \tag{3}$$

with $\alpha(t_n; \psi_k) = e^{-\beta_k t_n - \gamma_k t_n^2}$ denoting the envelope function. Assuming that the classes of the signal components are unknown, a straightforward approach would be, having estimated all parameters, to conduct a hypothesis test in order to determine whether $\beta_k > 0$ and/or $\gamma_k > 0$, for each individual component. However, finding the solution to (2) generally requires a $3K$-dimensional search, i.e., over ω_k, β_k, and γ_k, as solving for r_k and φ_k may be done using ordinary least squares. As an alternative, one may apply sparse reconstruction techniques by constructing a dictionary over all candidate parameters, reminiscent to the method in [9], but such a scheme quickly becomes practically infeasible due to the size of the problem for any reasonably fine grid for the different parameters. Motivated by mismatched estimation, we here instead formulate a sequential estimation scheme avoiding these difficulties, while still yielding an efficient estimator.

3 Mismatched estimation

Consider an observation of the signal in (1) for the special case when $K = 1$, collected in the vector $\mathbf{y} = \begin{bmatrix} y_0 & \dots & y_{N-1} \end{bmatrix}^T$. The corresponding probability density function (pdf) is then given by

$$p(\mathbf{y}; \psi, \alpha) = \frac{1}{(\pi\sigma^2)^N} \exp\left\{ -\frac{1}{\sigma^2} \sum_{n=0}^{N-1} |y_n - \mu(t_n; \psi, \alpha)|^2 \right\}, \tag{4}$$

where σ^2 is the variance of the additive noise. It is worth noting that the pdf is parametrized by the envelope function α, allowing for expressing the difference between the pdf for sinusoidal, Lorentzian, and Voigt models solely in terms of α. Specifically, we are concerned with the implication of parameter estimation in the case when α is replaced with a misspecified version $\breve{\alpha}$. We do this by considering the following definition:

Definition 1 (Pseudo-true parameter [15])**.** Consider a signal sample with pdf p, parametrized by the parameter vector ψ. For another parametric pdf, $\breve{p}$, parametrized by ϑ, the pseudo-true parameter, ϑ_0, is defined as

$$\vartheta_0 = \arg\min_{\vartheta} \; -\mathbb{E}_p \left(\log \breve{p}(\mathbf{y}; \vartheta) \right).$$

Thus, ϑ_0 minimizes the Kullback-Leibler divergence between the assumed and true signal models. Interestingly, and of importance for the problem of mismatched estimation, the mismatched maximum likelihood estimator, i.e.,

$$\hat{\vartheta}_{MLE} = \arg\max_{\vartheta} \log \breve{p}(\mathbf{y}; \vartheta),$$

where $\mathbf{y}$ is sampled from p, converges, under some regularity conditions, to the pseudo-true parameter ϑ_0 as the signal-to-noise ratio (SNR), or sample size, depending on the application, increases [15]. Herein, we are interested in estimating the parameters of purely sinusoidal and Lorentzian models when the actual measured signal may be Lorentzian or Voigt, respectively. The following two propositions detail the expected behavior in the one-component case.

Proposition 1. *The pseudo-true parameter vector* $\vartheta_0 = [r_0, \omega_0, \varphi_0]$ *corresponding to a model with* $\breve{\alpha}(t_n; \vartheta) \equiv 1$*, when* $\mathbf{y}$ *is sampled from* (4)*, is given by*

$$r_0 = \frac{r}{N} \sum_{n=0}^{N-1} \alpha(t_n; \psi) \quad , \quad \omega_0 = \omega \quad , \quad \varphi_0 = \varphi.$$

Proof. See the appendix. □

Thus, the estimate of the frequency and phase parameters are asymptotically unbiased, whereas the amplitude r will be underestimated. Further, for the mismatched estimation of a Lorentzian component, the following proposition holds.

Proposition 2. *Consider estimating the parameters* $\vartheta_0 = [r_0, \varphi_0, \omega_0, \beta_0]$ *corresponding to a model with* $\breve{\alpha}(t_n; \vartheta) = e^{-\beta t_n}$*, when* $\mathbf{y}$ *is sampled from* (4)*. Then,* $\omega_0 = \omega$ *and* $\varphi_0 = \varphi$*. Furthermore,*

$$r_0 = r \frac{\sum_{n=0}^{N-1} \alpha(t_n; \psi) e^{-\beta_0 t_n}}{\sum_{n=0}^{N-1} e^{-2\beta_0 t_n}},$$

where the pseudo-true decay parameter satisfies

$$\beta_0 \in \left(\beta, \beta + \gamma\,(t_{N-1} + t_{N-2})\right]$$

for any $\gamma > 0$.

Proof. See the appendix. □

It should be noted that the estimate of the frequency and phase parameters are thus unbiased, even if using an erroneous model, whereas the estimate of the linear decay parameter will incur a strictly positive bias. This result may be used as a bound when forming a final estimate of the parameter β. It should be stressed that the results of Propositions 1 and 2 are only exact for the single-component case. However, they will hold approximately for multi-component signals that do not contain components that are too closely spaced in frequency.

In the following section, we address the problem of how one may proceed to, in a computationally efficient manner, i.e., without having to form estimates of parameters of more complex models, determine whether a fitted model is sufficient for describing the measured data.

4 Spectrum test

As noted earlier, a Lorentzian model may be obtained as a special case of a Voigt model by setting $\gamma = 0$, and a sinusoidal model may be obtained by setting $\beta = \gamma = 0$. Thus, one may discern between the different models using standard hypothesis tests, i.e., by considering the statistical significance of the signal parameters [13]. However, such a procedure then requires estimating the full set of signal parameters. In order to avoid fitting unnecessarily complex models, we propose to exploit the difference in spectral properties of the three considered models. As detailed in Propositions 1 and 2, we expect that, for signals that do not contain too closely spaced components, the estimates of the frequency and phase parameters will be unbiased. Thus, components estimated under mismatched model assumptions will have spectra with the same modes as the actual signal, but with erroneous shapes. Also, in the residual spectra, we expect the power to be concentrated in a neighborhood of the estimated frequencies. With this observation, we propose to decide whether a fitted model is sufficient by considering the whiteness of the residual spectrum using the following proposition (see, e.g., [16]).

Proposition 3. *Under the null-hypothesis that the signal template coincides with the measurement model, it holds that the periodogram spectral estimate is distributed according to*

$$\frac{2\hat{\Phi}^{\mathrm{per}}(\omega)}{\sigma^2} \sim \chi^2(2),$$

for $\omega = k/2\pi N$, with $k = 1, 2, \ldots, N-1$ and σ^2 denoting the additive noise variance under the null hypothesis. Further, the test statistic

$$\xi(\hat{\omega}) \triangleq \frac{\frac{1}{|\mathcal{I}(\hat{\omega})|}\sum_{k\in\mathcal{I}(\hat{\omega})}\hat{\Phi}^{\text{per}}(\omega_k)}{\frac{1}{N-|\mathcal{I}(\hat{\omega})|}\sum_{k\notin\mathcal{I}(\hat{\omega})}\hat{\Phi}^{\text{per}}(\omega_k)},$$

where $\mathcal{I}(\hat{\omega})$ is a set of indices such that $\left\{\omega_k = \frac{k}{N2\pi} \mid k \in \mathcal{I}(\hat{\omega})\right\}$ is distributed according to

$$\xi(\hat{\omega}) \sim \text{F}\left(|\mathcal{I}(\hat{\omega})|, N - |\mathcal{I}(\hat{\omega})|\right),$$

where $\text{F}(d_1, d_2)$ denotes an F-distribution with (d_1, d_2) degrees of freedom.

Letting $\mathcal{I}(\hat{\omega})$ be a set of indices corresponding to a neighborhood of the estimated frequency $\hat{\omega}$ (e.g., ± 5 grid points), one may thus use Proposition 3 for detecting residual spectral power by comparing the test statistic $\xi(\hat{\omega})$ to quantiles of the F-distribution. In the numerical examples, the 95%-quantile is used. The benefit of this approach, as compared to standard hypothesis testing, is that one does not have to estimate the parameters of the more complex model in order to decide whether it is needed or not; all that is required is a candidate, potentially mismatched, the model, and an estimated spectrum. For simplicity, we here consider the periodogram estimate, although one could envision more sophisticated alternatives, although the distribution of the test statistic may in such case be different. In practice, for the case of multi-component signals, the set of spectral points used to estimate the noise power may be taken to be the complement of the union of all neighborhoods $\mathcal{I}(\hat{\omega})$, as opposed to the complement of the single neighborhood $\mathcal{I}(\hat{\omega})$ in Proposition 3.

5 Proposed algorithm

Based on the results in the previous section, we propose to address the estimation problem using the following procedure[3]

1. Fit a purely sinusoidal model to the signal of interest.

[3]The further steps are here only taken if the currently used model at each step is deemed to have been insufficient.

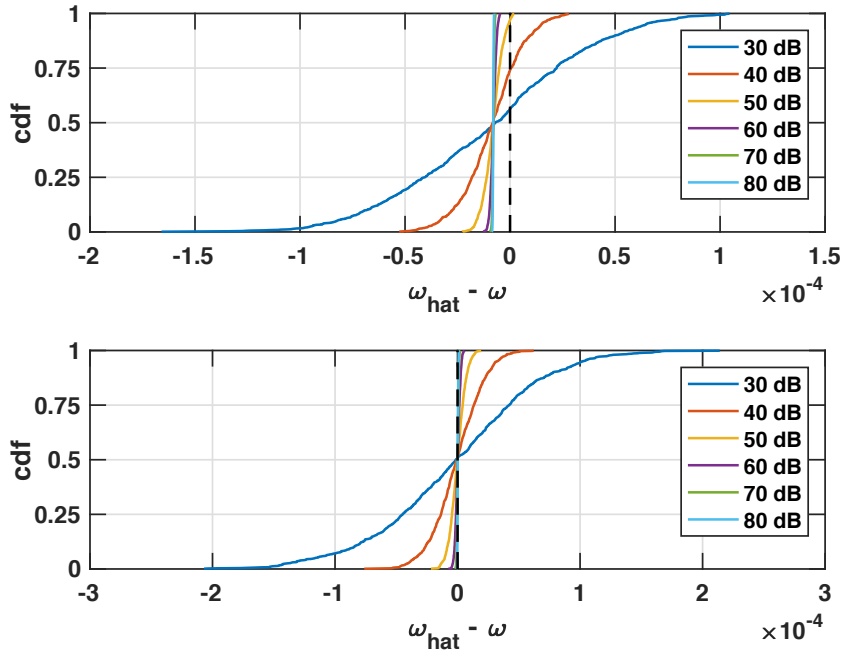

Figure 1: Empirical distribution functions for the error of initial frequency estimates. Top: Lorentzian component. Note that the Lorentzian and sine components are closely spaced. Bottom: Voigt component. Note that the Voigt component is well separated for the other two signal components.

2. Identify the damped components using Proposition 3. Add the identified components to signal residual.

3. Fit a Lorentzian model to the signal residual.

4. Identify the non-linearly damped components using Proposition 3. Add these to the signal residual.

5. Fit a Voigt model to the signal residual.

For the first and third steps, we propose to use the estimators SURE-IR [17] and D-SURE [11]. These estimators are approximate ML estimators tailored for sinusoidal and Lorentzian models, respectively, and both include techniques for arriving at sparse estimates not requiring *a priori* model order knowledge. Also,

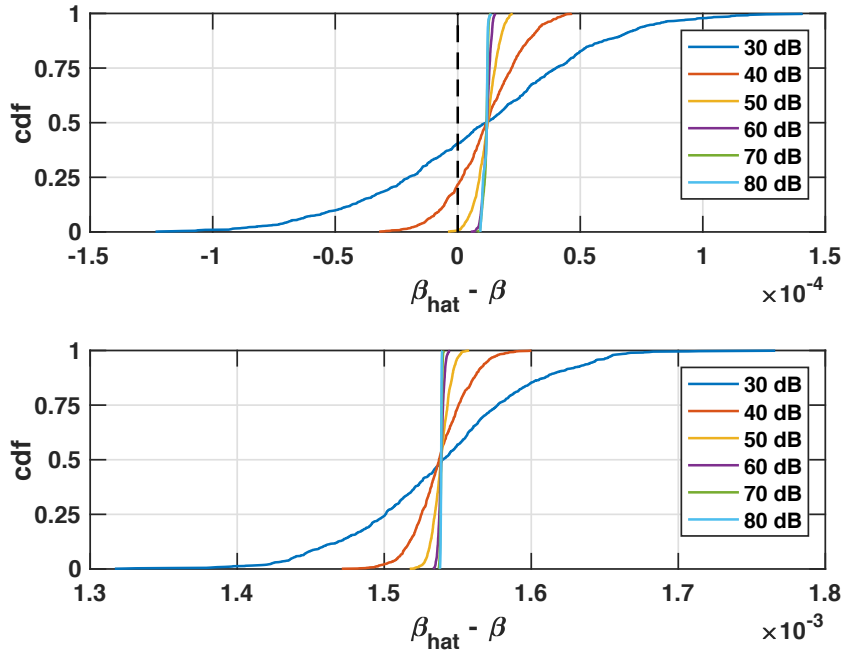

Figure 2: Empirical distribution functions for the error of initial linear decay estimates. Top: Lorentzian component. Bottom: Voigt component. Note that the bound for the estimation bias is 4×10^{-3}.

by not relying on a fixed gridding of the parameter space, the need for large signal dictionaries is alleviated, allowing for rapid implementation. In practice, the complexity of these estimators are $\mathcal{O}(\hat{K}^3)$, where $\hat{K}$ is the number of signal components identified by the estimators. Clearly, steps 1 and 3 may also be implemented using other estimators, such as, e.g., root-MUSIC [18] and its extensions [6, 7], although such methods in general require accurate *a priori* model order knowledge. In any case, in order to utilize the results from Propositions 1 and 2, the chosen estimators should approximate the MLE for the respective models.

As noted earlier, one could, as an alternative to the proposed sequential estimation, instead form a full dictionary that spans the whole parameter space and thus contains a (large) set of signal candidates and estimate the parameters by solving a sparsity enforcing optimization problem, e.g. Lasso [10, 19]. However, this approach has the major drawback of requiring a large number of candidate

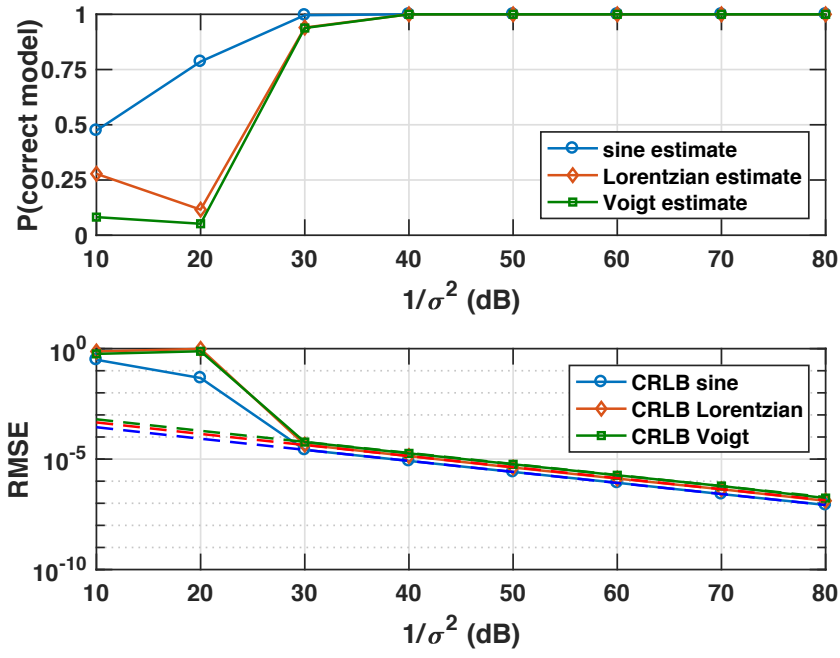

Figure 3: Top: Empirical probability of correctly classifying the three signal components. Bottom: Root-MSE for the frequency parameters, ω_k.

signal components in order to form reliable estimates. For instance, if one is to form the dictionary with P grid point for each parameter, the resulting dictionary would be of the size $N \times P^3$. To yield a satisfactory estimation precision for each parameter, P generally has to be quite large, often in the range of $10^2 - 10^4$. Solving problems on this scales requires significant computational resources and ignores available information.

In contrast, steps 1-4 above allows for an efficient estimation of the parameters of sinusoidal and Lorentzian components, as well as identifying Voigt components. Using Proposition 2, one may then use a non-linear search to efficiently estimate the Voigt parameters, as the proposition allows a limit on the relevant search space. As a final step, a local refinement search in a limited neighborhood of the estimated parameters may be performed. It should be noted that the proposed scheme allows for the dimension of the estimate to be kept to a minimum in every step by sequentially identifying and estimating the signal components.

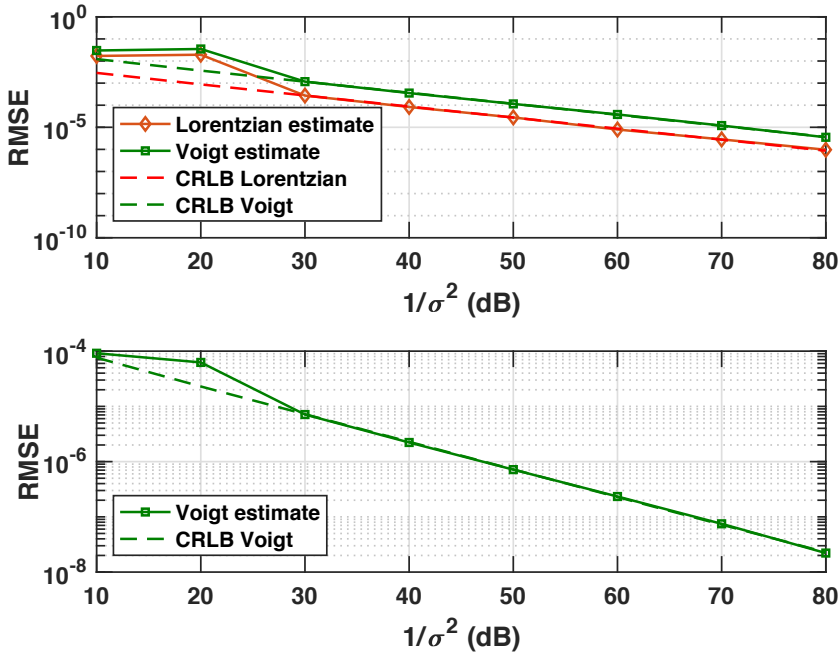

Figure 4: Top: Root-MSE for the linear decay parameters, β_k. Bottom: Root-MSE for the quadratic decay parameter, γ.

6 Numerical illustration

We consider a signal consisting of a pure sinusoidal, a Lorentzian, and a Voigt component, measured at $t_n = n$, for $n = 0, 1, \ldots, N-1$, with $N = 200$. The signal parameters are $(\omega_1, \omega_2, \omega_3) = (0.7, 0.5, 1.5)$, $(\beta_2, \beta_3) = (\frac{1}{200}, \frac{1}{150})$, and $\gamma_3 = 10^{-5}$, where the indices 1, 2, and 3 correspond to the sinusoidal, Lorentzian, and Voigt components, respectively, and $r_1 = r_2 = r_3 = 1$, with the phases drawn uniformly on $[0, 2\pi)$. We add circularly symmetric white Gaussian noise, and attempt to recover the signal parameters using the proposed estimation method. We repeat this procedure in 500 Monte Carlo simulations for different levels of noise.[4] Figure 1 displays the empirical cumulative distribution function for the error of the frequency estimates for the Lorentzian and Voigt components

[4]We stress that the estimator is not provided with oracle knowledge of the number of signal components, nor their specific signal category.

obtained in Step 1, i.e., when fitting a purely sinusoidal model. There is a small negative estimation bias for the Lorentzian component, caused by its proximity to the sine component, whereas the estimate for the Voigt component, that is well-separated from the other two components, is unbiased. Figure 2 displays corresponding results for the linear decay parameters obtained in Step 3, i.e., when fitting Lorentzian components. The estimate corresponding to the Voigt component has a substantial bias, given the scale of the parameters, although it is within the upper bound 4×10^{-3} predicted by Proposition 2. The empirical probability of correct model classification, i.e., the probability of determining $K = 3$ as well as correctly classifying the model type for each component, is displayed in the top panel of Figure 3. As can be seen, the probability approaches 1, for all three components, as the noise variance falls below 10^{-3}. Considering the simulations in which the components were correctly classified, the obtained root-MSE, being refined with a local non-linear search, for the signal parameters is shown in the bottom panels of Figure 3 and 4, as compared with the corresponding CRLB.

7 Appendix

7.1 Proof of Proposition 1

Proof. Excluding constant terms as well as positive scalings, expectation is given by

$$2r_0 r \sum_{n=0}^{N-1} \breve{\alpha}(t_n) \mathfrak{Re}\left(e^{i(\omega_0-\omega)t+i(\varphi_0-\varphi)}\right) - Nr_0^2.$$

Here, $\mathfrak{Re}\left(e^{i(\omega_0-\omega)t+i(\varphi_0-\varphi)}\right) \leq 1$, with equality being achieved for all t_n if and only if $\omega_0 = \omega$ and $\varphi_0 = \varphi + k2\pi$, for $k \in \mathbb{Z}$. Substituting this in and minimizing the negative of the resulting expression with respect to r_0 yields the above stated result. □

7.2 Proof of Proposition 2

Proof. The values of ω_0, φ_0, and r_0 may be obtained using the same reasoning as in the proof of Proposition 1. To find β_0, one needs to maximize

$$\Lambda(\beta_0) = \frac{\left(\sum_{n=0}^{N-1} e^{-\beta t_n - \gamma t_n^2 - \beta_0 t_n}\right)^2}{\sum_{n=0}^{N-1} e^{-2\beta_0 t_n}}.$$

Differentiation with respect to (w.r.t) β_0 yields

$$\frac{\partial \Lambda(\beta_0)}{\partial \beta_0} = \frac{\lambda_1(\beta_0) + \lambda_2(\beta_0)}{q(\beta_0)},$$

where q is a strictly positive function and

$$\lambda_1(\beta_0) = -2\left(\sum_{n=0}^{N-1} e^{-\beta t_n - \gamma t_n^2 - \beta_0 t_n}\right)\left(\sum_{n=0}^{N-1} t_n e^{-\beta t_n - \gamma t_n^2 - \beta_0 t_n}\right)\left(\sum_{n=0}^{N-1} e^{-2\beta_0 t_n}\right)$$

and

$$\lambda_2(\beta_0) = 2\left(\sum_{n=0}^{N-1} e^{-\beta t_n - \gamma t_n^2 - \beta_0 t_n}\right)^2\left(\sum_{n=0}^{N-1} t_n e^{-2\beta_0 t_n}\right).$$

Excluding common positive terms, one may arrive at the function (note explicit dependence on γ)

$$\begin{aligned}\Psi(\beta_0, \gamma) = &\left(\sum_{n=0}^{N-1} e^{-\beta t_n - \gamma t_n^2 - \beta_0 t_n}\right)\left(\sum_{n=0}^{N-1} t_n e^{-2\beta_0 t_n}\right) \\ &- \left(\sum_{n=0}^{N-1} t_n e^{-\beta t_n - \gamma t_n^2 - \beta_0 t_n}\right)\left(\sum_{n=0}^{N-1} e^{-2\beta_0 t_n}\right) \\ = &\sum_{m=1}^{N-1}\sum_{n=0}^{m-1}(t_m - t_n)e^{-\beta_0(t_m + t_n)} \\ &\cdot \left(e^{-\beta t_n - \gamma t_n^2 - \beta_0 t_m} - e^{-\beta t_m - \gamma t_m^2 - \beta_0 t_n}\right)\end{aligned}$$

which has the same zeros w.r.t. to β_0 and sign as the derivative of the original objective function Λ. It may be noted that $\Psi(\beta, 0) = 0$, i.e., the $\beta_0 = \beta$ for $\gamma = 0$. However, for $\gamma > 0$, we will show that $\beta_0 > \beta$ by showing that $\Psi(\tilde{\beta}, \gamma) > 0$ for any $\tilde{\beta} \in [0, \beta]$ and $\gamma > 0$. In fact, all terms of the last expression of $\Psi(\tilde{\beta}, \gamma)$ are positive for $\tilde{\beta} \in [0, \beta]$. To see this, note that $(t_m - t_n)e^{-\tilde{\beta}(t_m + t_n)} > 0$ for $t_m > t_n$. Thus, positivity is equivalent to

$$e^{-\beta t_n - \gamma t_n^2 - \tilde{\beta} t_m} > e^{-\beta t_m - \gamma t_m^2 - \tilde{\beta} t_n} \iff \frac{\tilde{\beta} - \beta}{\gamma} < t_m + t_n.$$

As $\tilde{\beta} - \beta \leq 0$ for $\tilde{\beta} \in [0, \beta]$, the inequality holds for all t_m and t_n. Thus, $\Psi(\tilde{\beta}, \gamma) > 0$ for all $\tilde{\beta} \in [0, \beta]$ as it is a sum of strictly positive terms. Using the same line of reasoning, we have that all terms of the sum are non-positive if $\tilde{\beta} \geq \beta + \gamma(t_{N-1} + t_{N-2})$, i.e., there exists a finite $\tilde{\beta}$ such that $\Psi(\tilde{\beta}, \gamma) \leq 0$. As $\Psi(\tilde{\beta}, \gamma)$ is continuous in $\tilde{\beta}$, we may conclude that $\Psi(\tilde{\beta}, \gamma) = 0$ for some $\tilde{\beta} \in \big(\beta, \beta + \gamma(t_{N-1} + t_{N-2})\big]$, which concludes the proof. □

7.3 A note on estimation performance

When estimating the parameters of a template model not coinciding with the model generating the measurements, standard bounds on estimation performance, such as the CRLB, are not applicable to the variance of estimators of the pseudo-true parameter. However, for estimators $\hat{\vartheta}$ of ϑ_0 that are unbiased under the true signal pdf p, a lower bound on the estimator variance is provided by the misspecified CRLB (MCRLB) [20], i.e.,

$$\mathbb{E}_p\left((\hat{\vartheta} - \vartheta_0)(\hat{\vartheta} - \vartheta_0)^T\right) \succeq \text{MCRLB}(\vartheta_0).$$

For the case of fitting a sinusoidal model to signal measurements sampled from (4), the MCRLB may be expressed in closed form solely in terms of the true parameter vector ψ. In fact, it may be shown that

$$\text{MCRLB}(\vartheta_0) = \frac{\sigma^2}{2}\begin{bmatrix} \frac{1}{N} & 0 & 0 \\ 0 & & \\ 0 & & \frac{1}{r^2}A(\psi)^{-1}BA(\psi)^{-1} \end{bmatrix},$$

where σ^2 is the variance of the additive noise and where

$$A(\psi) = \begin{bmatrix} \sum_{n=0}^{N-1} t_n^2\alpha(t_n) & \sum_{n=0}^{N-1} t_n\alpha(t_n) \\ \sum_{n=0}^{N-1} t_n\alpha(t_n) & \sum_{n=0}^{N-1} \alpha(t_n) \end{bmatrix}, B = \begin{bmatrix} \sum_{n=0}^{N-1} t_n^2 & \sum_{n=0}^{N-1} t_n \\ \sum_{n=0}^{N-1} t_n & N \end{bmatrix},$$

with $\alpha(t_n) \equiv \alpha(t_n; \psi)$. This reveals an interesting asymptotic property: assuming that the sampling times t_n constitute a strictly increasing sequence, the lower bounds on the variances for estimates of ω_0 and φ_0 diverge if $\beta > 0$ or $\gamma > 0$, as $N \to \infty$. To see this, note that as $\alpha(t)$ goes to zero exponentially fast as $t \to \infty$ for non-zero decay parameters, implying that $A(\psi)$, and thereby also $A(\psi)^{-1}$, converge to a finite positive definite matrix as $N \to \infty$. In contrast, the elements of B tend to infinity as $N \to \infty$, implying that $A(\psi)^{-1}BA(\psi)^{-1}$ diverges. Thus, for

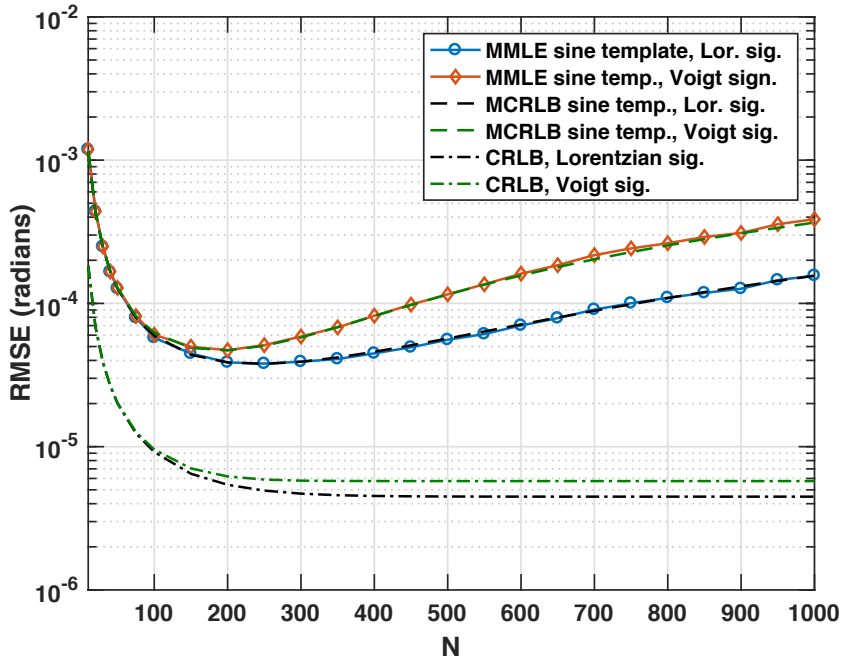

Figure 5: Root-MSE for the MMLE corresponding to a sinusoidal template applied to measurements from a Lorentzian and Voigt signal, respectively, for varying sample sizes N. Here, the noise variance is $\sigma^2 = 10^{-3}$.

large enough sample sizes N, the estimation performance will deteriorate as more samples are added. This is shown in Figure 5, displaying the (root) MCRLB for the pseudo-true frequency ω_0 corresponding to a sinusoidal template when the signal is either Lorentzian or Voigt. As may be noted, after initially decreasing, the bounds are strictly increasing for N larger than 250. Also shown is the root MSE obtained using the mismatched MLE (MMLE) in a Monte Carlo simulation study with 500 simulations for each sample size N. It may be noted that the MCRLB correctly predicts the performance of the MMLE. As reference, the corresponding CRLBs for the Lorentzian and Voigt models are also provided. As can be seen, the CRLBs converge to non-zero numbers as $N \to \infty$.

References

[1] L. Vanhamme, T. Sundin, P. van Hecke, and S. van Huffel, "MR spectroscopy quantitation: a review of time-domain methods," *NMR in Biomedicine*, vol. 14, no. 4, pp. 233–246, 2001.

[2] U. Dahlén, N. Rydén, and A. Jakobsson, "Damage Identification in Concrete using Impact Non-Linear Reverberation Spectroscopy," *NDT & E International*, vol. 75, pp. 15–25, 2015.

[3] M. Akke, R. Brüschweiler, and A. G. Palmer, "NMR order parameters and free energy: an analytical approach and its application to cooperative Ca2+ binding by calbindin-D(9K)," *J. Am. Chem. Soc.*, vol. 115, pp. 9832–9833, 1993.

[4] J. Higinbotham and I. Marshall, "NMR Lineshapes and Lineshape Fitting Procedures," *Annu. Rep. NMR Spectrosc.*, vol. 43, pp. 59–120, 2001.

[5] I. Marshall, J. Higinbotham, S. Bruce, and A. Freise, "Use of Voigt lineshape for quantification of in vivo 1H spectra," *Magn. Res. in Medicine*, vol. 37, pp. 651–657, 2005.

[6] S. Van Huffel, H. Chen, C. Decanniere, and P. Van Hecke, "Algorithm for Time-Domain NMR Data Fitting based on Total Least Squares," *J. Magn. Reson.*, vol. 110, pp. 228–237, 1994.

[7] F. K. W. Chan, H. C. So, and W. Sun, "Subspace approach for two-dimensional parameter estimation of multiple damped sinusoids," *Signal Process.*, vol. 92, pp. 2172 – 2179, 2012.

[8] J. Klintberg and T. McKelvey, "An Improved Method for Parametric Spectral Estimation," in *Proc. 44th IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, Brighton, UK, May 13-17 2019, pp. 5551–5555.

[9] S. Sahnoun, E. Djermoune, C. Soussen, and D. Brie, "Sparse multidimensional modal analysis using a multigrid dictionary refinement," *EURASIP J. Applied SP*, vol. 60, pp. 1–10, 2012.

[10] J. Swärd, S. I. Adalbjörnsson, and A. Jakobsson, "High Resolution Sparse Estimation of Exponentially Decaying N-dimensional Signals," *Elsevier Signal Processing*, vol. 128, pp. 309–317, Nov 2016.

[11] M. Juhlin, F. Elvander, J. Swärd, and A. Jakobsson, "Fast Grid-less Estimation of Damped Modes," in *2018 Int. Symp. on Intel. Sig. Proc. and Com. Syst.*, Ishigaki Island, Okinawa, Japan, November 2018.

[12] S. D. Bruce, J. Higinbotham, I. Marshall, and P. H. Beswick, "An Analytical Derivation of a Popular Approximation of the Voigt Function for Quantification of NMR Spectra," *J. Magn. Reson.*, vol. 142, no. 1, pp. 57 – 63, 2000.

[13] S. I. Adalbjörnsson and A. Jakobsson, "Relax-Based Estimation of Voigt Lineshapes," in *18th European Signal Processing Conference, EUSIPCO 2010*, Aalborg, Denmark, 2010, pp. 1053–1057.

[14] P. Stoica and Y. Selén, "Model-order Selection — A Review of Information Criterion Rules," *IEEE Signal Process. Mag.*, vol. 21, no. 4, pp. 36–47, July 2004.

[15] S. Fortunati, F. Gini, M. S. Greco, and C. D. Richmond, "Performance bounds for parameter estimation under misspecified models: Fundamental findings and applications," *IEEE Signal Process. Mag.*, vol. 34, no. 6, pp. 142–157, Nov 2017.

[16] A. Jakobsson, *Time Series Analysis and Signal Modeling*, Studentlitteratur AB, Sweden, 3 edition, 2019.

[17] J. Fang, F. Wang, Y. Shen, H. Li, and R. S. Blum, "Super-Resolution Compressed Sensing for Line Spectral Estimation: An Iterative Reweighted Approach," *IEEE Trans. Signal Process.*, vol. 64, no. 18, pp. 4649–4662, September 2016.

[18] B. D. Rao and K. V. S. Hari, "Performance Analysis of Root-MUSIC," *IEEE Trans. Acoust., Speech, Language Process.*, vol. 37, no. 12, pp. 1939–1949, December 1989.

[19] T. Hastie, R. Tibshirani, and M. Wainwright, *Statistical Learning with Sparsity: The Lasso and Generalizations*, Chapman and Hall/CRC, 2015.

[20] C. D. Richmond and L. L. Horowitz, "Parameter Bounds on Estimation Accuracy Under Model Misspecification," *IEEE Trans. Signal. Process*, vol. 63, no. 9, pp. 2263–2278, 2015.

E

Paper E

Designing Sampling Schemes for Multi-Dimensional Data

Johan Swärd, Filip Elvander, and Andreas Jakobsson

Centre for Mathematical Sciences, Lund University, Lund, Sweden

Abstract

In this work, we propose a method for determining a non-uniform sampling scheme for multi-dimensional signals by solving a convex optimization problem reminiscent of the sensor selection problem. The resulting sampling scheme minimizes the sum of the Cramér-Rao lower bounds for the parameters of interest, given a desired number of sampling points. The proposed framework allows for selecting an arbitrary subset of the parameters detailing the model, as well as weighing the importance of the different parameters. Also presented is a scheme for incorporating any imprecise *a priori* knowledge of the locations of the parameters, as well as defining estimation performance bounds for the parameters of interest. Numerical examples illustrate the efficiency of the proposed scheme.

Key words: Sampling schemes, Cramér-Rao lower bound, convex optimization

1 Introduction

Determining how to suitably sample a signal is an important problem in many signal processing applications, such as sensor positioning and selection in network monitoring [1,2], localization and tracking [3], magnetic resonance imaging (MRI) [4], graph signal processing [5, 6], and selecting the temporal sampling [7]. In general, these problems can be viewed as sampling a multi-dimensional space containing partly known signal components. For high-dimensional data, it quickly becomes infeasible to sample the space uniformly, especially, in areas such as nuclear magnetic resonance (NMR) spectroscopy when examining living cells, which have limited lifetimes. For example, a recent study of 4-D NMR measurements that would have taken about 2.5 years to perform using regular sampling was shown to be possible to construct in merely 90 hours using a non-uniform sampling scheme [8]. This has caused an interest in formulating sampling schemes for NMR signals, allowing for notable improvements [7,9–13].

Among the developed schemes are some exploiting a compressive sensing framework, allowing for an accurate signal reconstruction using fewer samples than the Nyquist-Shannon sampling theorem necessitates for uniformly sampled signals (see, e.g., [11, 12, 14, 15]). However, the developed schemes typically do not optimize the sampling scheme with respect to the expected signals, even though these are often fairly well known. In this work, we strive to exploit this knowledge in order to design a sampling scheme that would allow for an optimal estimation accuracy given the assumed prior knowledge.

There are many related problems to the herein studied sampling scheme problem. In [16], the problem of how to optimally measure a signal in problems related to propagating wave fields was studied. More specifically, the authors studied how to best recover the input wave field from noisy measurements of the output field given that each measurement is associated with a cost, where the selected cost was set higher for measurement devices with better resolution. The results were presented as trade-off curves between the error of estimation and the total cost budget. In [17], a framework for joint hypothesis testing and estimation using a minimal sample size was developed. The proposed framework guarantees, under a Bayesian setup, that the overall detection and estimation performance, given the minimization of the sample size, is the best possible. In [18], the optimal placement of phasor measurement units on power grids was studied. Other works have been studying problems related to sampling in random fields [19, 20] and wireless sensor networks [21]. A notable example of the latter category is [21], where

the problem of target tracking in wireless sensor networks is studied. The sensors with the most information are found by utilizing a proposed probabilistic sensor management scheme based on the compressed sensing framework. This scheme is determined based on the probability of transmission at each node, found by maximizing the trace of the Fisher information matrix (FIM). Using this approach, sensors with less information can be discarded, implying that fewer sensors need to communicate, thus leading to energy savings.

Lately, for the related problem of optimal sensor placement, there has been several methods proposed in which the combinatorial problem of selecting a subset of sensors is relaxed using convex optimization. In [22], the authors consider the case when signal measurements are linear in the unknown parameters and propose a sensor selection scheme based on solving a convex optimization problem inspired by the determinant criterion (D-optimality) of experimental design [23]. This work was then further developed in [2, 24–27], wherein the authors consider non-linear measurement equations, as well as replacing D-optimality with the average variance criterion (A-optimality) as a performance measure. Specifically, as A-optimality can be interpreted as the sum of the diagonal elements of the Cramér-Rao lower bound (CRLB) for the signal parameters, the problem was formulated as to minimize the number of required sensors subject to an upper bound on the resulting diagonal sum of the CRLB. Assuming that the bound is tight, the method thus finds a sparse set of sensors, i.e., activates a few out of a set of candidate sensors, while keeping the variance of the estimated parameters below a fixed level.

In this paper, we expand on this idea, proposing a method for finding a suitable sampling scheme in order to estimate the parameters for signal models where, in general, the signal measurements are non-linear functions of the unknown parameters. By taking the available prior information of the signal into consideration, we propose a sampling scheme that is found by solving a convex optimization problem that guarantees a bound on the worst case CRLB. The sampling pattern is selected via a variable vector, corresponding to the available sample positions, which is penalized using the ℓ_1-norm, resulting in a sampling scheme that is limited in the number of samples. Furthermore, we reformulate the optimization problem into a semidefinite programming (SDP) problem that allows for more flexibility and can be used for adding additional constraints on the optimization. In general, when estimating a set of parameters, it might be that the scale of the parameters, as well as the accuracy with which they can be es-

timated, are significantly different. Also, some of the unknown parameters might be of greater interest than the others; again, using NMR as an example, the signal decay is often of more interest than the signal frequencies, the latter often being relatively well known for a given substance, whereas the former measures the sought interactions. We here propose to use a weighting scheme in order to allow for a relative balancing of the variances of the different parameters, allowing for designing sampling schemes specifically tailored to yield good estimation accuracy for the parameters of interest.

In some applications, one may assume some prior knowledge of the signal of interest, such as, for example, knowledge of the subspace where the signal parameters are to be found. Again using NMR as an illustrative example, the signals of interest consist of decaying modes, being well modeled as a sum of damped sinusoids. These modes are, as noted, often well known in frequency, at least within some reasonably well defined frequency band, whereas the uncertainty of, and the interest in, the signal decays is often more significant. Typically, the problem of interest is thus to specify the damping parameter as accurately as possible using as few samples as possible. To allow for this case, we herein propose using a gridding of the parameter space in order to guarantee performance within certain bounds, allowing for uncertainty in the parameters.

To summarize, the main components of this paper are:

- The development of a unified framework for incorporating strong prior knowledge when designing a sampling scheme, such that one may determine which samples that are most appropriate to select for a signal that may contain the specified components. It should be stressed that any assumed component does not actually need to be present in the signal; the scheme will select samples that are sufficient to estimate signal components at the considered locations well, but these may be absent in an actual measurement.

- The method allows samples to be selected such that some parameters are deemed more important, i.e., of higher interest for the application, than others, and the sample scheme is prioritized accordingly. For instance, one is often more interested in the decay rate of the signal components than their actual frequencies (which are commonly well known), and may then focus the sampling scheme to allow for accurate estimation of these parameters.

- The method allows for uncertainty in the assumed model, allowing a user to indicate regions of interest wherein parameters are assumed to lie. This allows the proposed method to design sampling patterns for signals that are only partly known, relaxing the assumption of the signal modelled being known in detail *a priori*.

- The proposed method is valid for high-dimensional data sets, with the achievable gain growing (rapidly) with the growing dimensionality of the problem.

To illustrate the proposed method, we consider multi-dimensional NMR signals. However, the proposed framework is not limited to this case. Indeed, it is applicable to any signal model having a parametric description that allows for the existence of a Fisher information matrix. Although, we show that the proposed method yields Cramér-Rao lower bounds that are attainable, it should be stressed that this paper is not concerned with parameter estimation as such; any efficient estimator allowing for non-uniformly sampled data may be applied.

This paper is organized as follows. In Section 2, we introduce the problem statement and derive the proposed optimization problem. In Section 3, we present extensive numerical simulations and results that validates our proposed method. Finally, in Section 4, we conclude upon our work.

2 Problem statement and proposed sampling scheme

Consider a measured signal $y(\mathbf{t}_n)$, defined on a D-dimensional space with N potential D-dimensional sampling points, $\mathbf{t}_n$, $n = 1, 2, \ldots, N$. It is assumed that the probability density function (pdf) of $y(\mathbf{t}_n)$, here denoted by $p\left(y(\mathbf{t}_n); \boldsymbol{\vartheta}\right)$, is parametrized by the parameter vector $\boldsymbol{\vartheta} \in \mathbb{R}^P$ and that two samples $y(\mathbf{t}_n)$ and $y(\mathbf{t}_m)$ are independent if $\mathbf{t}_n \neq \mathbf{t}_m$. The FIM for sample $y(\mathbf{t}_n)$ may then be defined as

$$\mathbf{F}(\mathbf{t}_n; \boldsymbol{\vartheta}) = \mathbb{E}\left\{ \nabla_{\boldsymbol{\vartheta}} \log\left(p(y(\mathbf{t}_n); \boldsymbol{\vartheta})\right) \nabla_{\boldsymbol{\vartheta}} \log\left(p(y(\mathbf{t}_n); \boldsymbol{\vartheta})\right)^H \right\},$$

where $\mathbb{E}\{\cdot\}$, $\nabla_{\boldsymbol{\vartheta}}$, and $(\cdot)^H$ denote the statistical expectation, the gradient with respect to $\boldsymbol{\vartheta}$, and the conjugate transpose, respectively. The herein proposed sampling scheme is designed such that it is optimal in the sense of either minimizing the CRLB of the parameters of interest, given that M of the N potential

uniform samples are used, or conversely, minimizing the number of samples used given a desired upper bound on the CRLB of the parameters. It is worth noting that as the potential signal samples are assumed to be independent, for any set of sample indices $\mathbf{\Omega}$, it holds that

$$\sum_{n\in\mathbf{\Omega}} \mathbf{F}(\mathbf{t}_n;\boldsymbol{\vartheta})$$

is the corresponding FIM using this sample scheme. Let the N-dimensional vector $\mathbf{w}$ denote an indicator for the possible sampling points in the D-dimensional sampling space, such that if the nth index, w_n, is set to one, this sampling point is used, whereas if it is set to zero, it is not. Reminiscent of the case of optimal sensor selection, the resulting sampling design problem may then be formulated as (see also [24])

$$\begin{aligned} \underset{\mathbf{w}}{\text{minimize}} \quad & \text{tr}\left(\left(\sum_{n=1}^{N} w_n \mathbf{F}(\mathbf{t}_n;\boldsymbol{\vartheta})\right)^{-1}\right) \\ \text{subject to} \quad & \|\mathbf{w}\|_1 \leq \gamma \\ & w_n \in \{0,1\}\ ,\ n=1,2,\ldots,N, \end{aligned} \tag{1}$$

where $\gamma > 0$ and $\text{tr}(\cdot)$ denotes the trace operator. The choice of objective function is related to the so-called A-optimality criterion from design of experiments [23] as the trace of the inverse FIM corresponds to the sum of the CRLBs of the signal parameters in $\boldsymbol{\vartheta}$. Here, the parameter γ constitutes an upper bound on the ℓ_1-norm of the sample selection vector. The sampling design scheme (1) is not convex due to the restriction that w_n, for $n = 1,\ldots,N$, is defined over a non-convex set. A convex approximation to this problem may be found by relaxing the binary constraint and instead allowing w_n to take any value in the range $[0, 1]$ (see, e.g., [25]), resulting in

$$\begin{aligned} \underset{\mathbf{w}}{\text{minimize}} \quad & \text{tr}\left(\left(\sum_{n=1}^{N} w_n \mathbf{F}(\mathbf{t}_n;\boldsymbol{\vartheta})\right)^{-1}\right) \\ \text{subject to} \quad & \mathbf{1}^T\mathbf{w} \leq \gamma \\ & w_n \in [0,1]\ ,\ n=1,2,\ldots,N, \end{aligned} \tag{2}$$

where $\mathbf{1}$ is a vector of ones of appropriate dimension. It should be noted that we can here replace $||\mathbf{w}||_1$ with simply $\mathbf{1}^T\mathbf{w}$, since each element in $\mathbf{w}$ is equal

to or greater than zero. Given a solution $\hat{\mathbf{w}}$ to (2), we define the FIM for the corresponding sampling pattern as

$$\mathcal{I}(\hat{\mathbf{w}};\boldsymbol{\vartheta}) = \sum_{\ell\in\boldsymbol{\Omega}} \mathbf{F}(\mathbf{t}_\ell;\boldsymbol{\vartheta}), \quad \boldsymbol{\Omega} = \{\ell \mid \hat{w}_\ell > \xi\}\,, \tag{3}$$

where $\xi \geq 0$ is a threshold determining whether a sample weight $\hat{w}_\ell$ should be rounded toward one or zero, i.e., whether the sampling point should be included or not. This formulation allows for the minimization of the sum of the CRLBs given an upper bound on the number of samples used. Note that the problem could alternatively be formulated as minimizing the number of sampling points given an upper bound on the sum of the CRLBs. It is also worth noting that the signal to noise ratio (SNR) will not affect the sampling scheme as such, as the noise power will only affect the optimisation problem by scaling the cost function. Obviously, the SNR will result the performance of any estimator subsequently used to estimate the signal parameters using the selected samples, but thus not the actual selection of these samples.

However, the sampling design in (2) does not allow for the case when one is primarily interested in a subset of the available parameters. Neither does the formulation take into account that the different parameters might have significantly different variances. For example, for a sum of damped sinusoids, the trace constraint in (2) will clearly be dominated by the CRLB for the amplitudes, as these are orders of magnitude larger than those for the frequencies, and the optimization will therefore put an emphasis on minimizing the CRLBs of the amplitude parameters. In order to allow for sampling schemes that put an emphasis on a selection of the parameters of interest, we recently proposed to introduce a weighting matrix, $\mathbf{A}(\boldsymbol{\vartheta})$, acting upon the FIM in [28]. Specifically, instead of minimizing the cost function using the FIM, we proposed to perform the minimization using weighted FIMs

$$\tilde{\mathbf{F}}(\mathbf{t}_n;\boldsymbol{\vartheta}) = \mathbf{A}(\boldsymbol{\vartheta})\mathbf{F}(\mathbf{t}_n;\boldsymbol{\vartheta})\mathbf{A}(\boldsymbol{\vartheta})^T,$$

i.e., performing a linear transformation of the variables and minimizing the sum of the CRLBs corresponding to the transformed parameters $\tilde{\boldsymbol{\vartheta}} = \mathbf{A}(\boldsymbol{\vartheta})\boldsymbol{\vartheta}$. However, although this formulation allows for shifting emphasis to the parameters of interest, it does not allow for complete disregard of nuisance parameters as $\mathbf{A}(\boldsymbol{\vartheta})$ has to be non-singular in order for the matrix inverse to be defined. In order to

allow for an arbitrary weighting, we note that the following useful identity holds for any invertible $P \times P$ matrix $\mathbf{B}$;

$$\operatorname{tr}\left(\mathbf{B}^{-1}\right) = \sum_{p=1}^{P} \mathbf{e}_p^T \mathbf{B}^{-1} \mathbf{e}_p,$$

where $\mathbf{e}_p$ denotes the pth canonical basis vector, i.e., a vector with all its elements equal to zero except the pth being equal to one. Furthermore, it is noted that for any positive definite matrix $\mathbf{B}$, scalar μ, and arbitrary vector $\mathbf{a}$, it follows from the Schur complement (see, e.g., [29]) that

$$\mu - \mathbf{a}^T \mathbf{B}^{-1} \mathbf{a} \geq 0 \iff \begin{bmatrix} \mathbf{B} & \mathbf{a} \\ \mathbf{a}^T & \mu \end{bmatrix} \succeq \mathbf{0},$$

where $\mathbf{X} \succeq \mathbf{0}$ indicates that the matrix $\mathbf{X}$ is positive semidefinite. Thus, it follows that

$$\underset{\mathbf{B} \succ 0}{\text{minimize}} \ \mathbf{a}^T \mathbf{B}^{-1} \mathbf{a}$$

and

$$\begin{aligned} \underset{\mu, \mathbf{B} \succ 0}{\text{minimize}} \quad & \mu \\ \text{subject to} \quad & \begin{bmatrix} \mathbf{B} & \mathbf{a} \\ \mathbf{a}^T & \mu \end{bmatrix} \succeq \mathbf{0} \end{aligned}$$

are minimized by the same matrix $\mathbf{B}$. Here, $\mathbf{B} \succ \mathbf{0}$ indicates that the matrix $\mathbf{B}$ is positive definite. This observation allows us to reformulate (2) as the semidefinite program (SDP) (cf. [2, 18])

$$\begin{aligned} \underset{\boldsymbol{\mu}, \mathbf{w}}{\text{minimize}} \quad & \sum_{p=1}^{P} \psi_p \mu_p \\ \text{subject to} \quad & \begin{bmatrix} \sum_{n=1}^{N} w_n \mathbf{F}(\mathbf{t}_n; \boldsymbol{\vartheta}) & \mathbf{e}_p \\ \mathbf{e}_p^T & \mu_p \end{bmatrix} \succeq \mathbf{0}, \ p = 1, \ldots, P \\ & \sum_{n=1}^{N} w_n \mathbf{F}(\mathbf{t}_n; \boldsymbol{\vartheta}) \succ \mathbf{0} \\ & \mathbf{1}^T \mathbf{w} \leq \gamma \quad , \quad w_n \in [0, 1], \ n = 1, \ldots, N, \end{aligned} \tag{4}$$

where ψ_p are weight parameters allowing for putting emphasis on different components of the vector $\boldsymbol{\vartheta}$. For example, if $\psi_q = 1$ and $\psi_p = 0$, $\forall p \neq q$, then the CRLB for the parameter ϑ_q will be the only one minimized, as μ_q precisely corresponds to this lower bound, whereas the CRLBs for the other parameters ϑ_p, $p \neq q$ will be disregarded. Similarly, for $\psi_p = 1$, $\forall p$, the problems (2) and (4) are equivalent.

Another benefit of this formulation is that it allows for a straightforward way of incorporating performance constraints in the minimization problem, such as if, for instance, there is some upper tolerance bound λ_p for the CRLB of parameter ϑ_p. This kind of performance specifications can then be incorporated in the minimization problem via linear inequality constraints according to

$$
\begin{aligned}
\underset{\boldsymbol{\mu},\mathbf{w}}{\text{minimize}} \quad & \sum_{p=1}^{P} \psi_p \mu_p \\
\text{subject to} \quad & \begin{bmatrix} \sum_{n=1}^{N} w_n \mathbf{F}(\mathbf{t}_n;\boldsymbol{\vartheta}) & \mathbf{e}_p \\ \mathbf{e}_p^T & \mu_p \end{bmatrix} \succeq \mathbf{0},\ p = 1,\ldots,P \\
& \sum_{n=1}^{N} w_n \mathbf{F}(\mathbf{t}_n;\boldsymbol{\vartheta}) \succ \mathbf{0}\ ,\ \mathbf{1}^T\mathbf{w} \leq \gamma \\
& w_n \in [0,1],\ n = 1,\ldots,N\ ,\ \mu_p \leq \lambda_p\ ,\ p = 1,\ldots,P.
\end{aligned}
\tag{5}
$$

Furthermore, one may not only be interested in designing a sampling scheme for a single parameter vector $\boldsymbol{\vartheta}$, but rather for a set of parameter vectors. For example, consider the case when the parameters in $\boldsymbol{\vartheta}$ are only partly known, such that one may assume that $\boldsymbol{\vartheta}$ instead lies in a set of possible parameters, Θ. In such cases, it may be desired to treat some of the parameters as known, whereas others are only partly known, within some set of uncertainty. To allow for this, as well as taking the weighting into account, we further generalize (5) such that the

sampling scheme is designed as

$$\begin{aligned}
\underset{\boldsymbol{\mu},\mathbf{w}}{\text{minimize}} \quad & \sum_{p=1}^{P} \psi_p \mu_p \\
\text{subject to} \quad & \begin{bmatrix} \sum_{n=1}^{N} w_n \mathbf{F}(\mathbf{t}_n;\vartheta) & \mathbf{e}_p \\ \mathbf{e}_p^T & \mu_p \end{bmatrix} \succeq \mathbf{0}, \ \forall \vartheta \in \Theta, \ p = 1,\ldots,P \\
& \sum_{n=1}^{N} w_n \mathbf{F}(\mathbf{t}_n;\vartheta) \succ \mathbf{0} \\
& \mathbf{1}^T \mathbf{w} \leq \gamma \quad , \quad w_n \in [0,1], \ n = 1,\ldots,N \\
& \mu_p \leq \lambda_p \ , \ p = 1,\ldots,P.
\end{aligned} \tag{6}$$

Using this formulation, the optimal μ_p will, assuming that $\psi_p > 0$, now correspond to a worst case CRLB for the pth component of ϑ, when $\vartheta \in \Theta$, i.e., for the obtained sampling sampling scheme

$$\mu_p = \arg\max_{\vartheta \in \Theta} \mathbf{e}_p^T \mathcal{I}(\hat{\mathbf{w}};\vartheta)^{-1} \mathbf{e}_p.$$

Thus, the solution to (6) is a sampling scheme minimizing the worst case CRLB for the parameters of interest if the parameter vector ϑ is known to be in the set Θ.

Further, one could also consider the case where there is some cost associated with changing sampling points in one of the dimensions. For instance, if one of the sampling dimensions corresponds to a certain setting of a machine, e.g., time delay or magnetic flow, it could be more costly to acquire many different sample points in this dimension. Illustrating this in the 2-D case, one could include such a cost in the optimization by forming the $N_1 \times N_2$ matrix $\mathbf{W}$ by reshaping the vector $\mathbf{w}$, and adding the constraints

$$\begin{aligned}
\left|\left|\mathbf{W}^T\right|\right|_{2,1} &= \sum_{n=1}^{N_1} \left|\left|\mathbf{W}_{(:,n)}\right|\right|_2 \leq \gamma_1 \\
||\mathbf{W}||_{2,1} &= \sum_{n=1}^{N_2} \left|\left|\mathbf{W}_{(n,:)}\right|\right|_2 \leq \gamma_2
\end{aligned}$$

to the optimization in (6). Here, γ_1 and γ_2 are tuning parameters that may be set according to the associated cost. This constraint can easily be omitted simply by setting $\gamma_1 = \gamma_2 = \infty$.

It is also worth noting that when relaxing (1) in favor for (2), we can no longer guarantee that the weights are exactly 0 or 1. In this case, as is noted in (3), we simple choose an appropriate threshold such that values above the threshold are deemed as ones, and the values below are deemed as zeros. However, a better approximation of (1) is found by using re-weighting. This may be done by first solving (6), yielding the estimated $\mathbf{w}^{(1)}$, where the superscript $(\cdot)^{(j)}$ denotes jth iteration. Then, (6) is solved again, but this time with $\frac{1}{w_n^{(1)}+\delta}$ as a scaling factor for each w_n, where δ is a small number added to the denominator to avoid numerical problems. This procedure can then repeated until convergence. The re-weighting is a better approximation of the ℓ_0-norm, and thus is more likely to produce weights with values close to zero or one. As we have empirically found that using re-weighting for the here studied examples offers only a marginal improvement, while significantly increasing the computational cost due to the iterative procedure, we have in our examples chosen to use the simpler thresholding approach.

The SDP formulations in equations (4)-(6) may be minimized using, e.g., off-the-shelf solvers such as SeDuMi [30] or SDPT3 [31]. Typically, these require $\mathcal{O}(J^3)$ operations per iteration, where J is the number of variables. Clearly, this implies that the solution will be computationally demanding for large potential sampling grids and high-dimensional data. However, as noted in the introduction, in many experimental applications, the time required to carry out an experiment is often several orders of magnitude greater than the time needed for determining the optimal sampling scheme, even using standard solvers not exploiting any inherent structure in the problem. As a result, the scheme allows for a substantial overall decrease of the experimental time, even if the forming of the sampling scheme is somewhat cumbersome, as this step only needs to be done once for a given experimental setup, whereas the resulting experiment would typically be done multiple times, with each measurement typically taking substantially longer than the design of the sampling scheme.

3 Numerical results

3.1 Illustration in 1-D

To illustrate the proposed sampling scheme, we consider the NMR signal model, as noted being formed as a sum of damped sinusoids (for ease of notation, we

initially focus on the 1-D case), such that

$$y(t_n) = \sum_{k=1}^{K} \alpha_k \exp\left(2i\pi f_k t_n - \beta_k t_n + i\varphi_k\right) + \varepsilon(t_n), \tag{7}$$

for $n = 1, \ldots, N$, where α_k, f_k, β_k, and φ_k are the amplitude, frequency, damping, and phase of the kth component, respectively, and where ε is an additive noise term, here assumed to be well modeled as a white circularly symmetric Gaussian noise[1] with variance σ^2, with t_n being the time at sample n. For simplicity, we consider uniformly sampled candidate sampling times, t_n. For the signal model in (7), the FIM corresponding to sampling time t_n is given by

$$\mathbf{F}(t_n;\boldsymbol{\vartheta}) = \frac{2}{\sigma^2}\begin{pmatrix} \mathbf{F}_{1,1}(t_n) & \mathbf{F}_{1,2}(t_n) & \ldots & \mathbf{F}_{1,K}(t_n) & 0 \\ \vdots & \vdots & \ldots & \vdots & \vdots \\ \mathbf{F}_{K,1}(t_n) & \mathbf{F}_{K,2}(t_n) & \ldots & \mathbf{F}_{K,K}(t_n) & 0 \\ 0 & 0 & \ldots & 0 & 1/\sigma^2 \end{pmatrix},$$

where

$$\mathbf{F}_{k,\ell}(t_n) = \begin{pmatrix} \frac{c_{k,\ell}(t_n)}{\alpha_k\alpha_\ell} & \frac{2\pi t_n s_{k,\ell}(t_n)}{\alpha_k} & \frac{-t_n c_{k,\ell}(t_n)}{\alpha_k} & \frac{s_{k,\ell}(t_n)}{\alpha_k} \\ \frac{-2\pi t_n s_{k,\ell}(t_n)}{\alpha_\ell} & (2\pi t_n)^2 c_{k,\ell}(t_n) & 2\pi t_n^2 s_{k,\ell}(t_n) & 2\pi t_n c_{k,\ell}(t_n) \\ \frac{-t_n c_{k,\ell}(t_n)}{\alpha_\ell} & -2\pi t_n^2 s_{k,\ell}(t_n) & t_n^2 c_{k,\ell}(t_n) & -t_n s_{k,\ell}(t_n) \\ \frac{-s_{k,\ell}(t_n)}{\alpha_\ell} & 2\pi t_n c_{k,\ell}(t_n) & t_n s_{k,\ell}(t_n) & c_{k,\ell}(t_n) \end{pmatrix},$$

with

$$\begin{aligned} c_{k,\ell}(t_n) &= \alpha_k\alpha_\ell e^{-(\beta_k+\beta_\ell)t_n}\cos\left(2\pi(f_k - f_\ell)t_n + (\varphi_k - \varphi_\ell)\right) \\ s_{k,\ell}(t_n) &= \alpha_k\alpha_\ell e^{-(\beta_k+\beta_\ell)t_n}\sin\left(2\pi(f_k - f_\ell)t_n + (\varphi_k - \varphi_\ell)\right). \end{aligned}$$

Here, the parameter vector $\boldsymbol{\vartheta}$ is defined as

$$\begin{aligned} \boldsymbol{\vartheta} &= \begin{bmatrix} \boldsymbol{\vartheta}_1^T & \ldots & \boldsymbol{\vartheta}_K^T \end{bmatrix}^T \\ \boldsymbol{\vartheta}_k &= \begin{bmatrix} \alpha_k & f_k & \beta_k & \varphi_k \end{bmatrix}^T. \end{aligned}$$

As an illustration, Figure 1 shows an example of sampling schemes found by

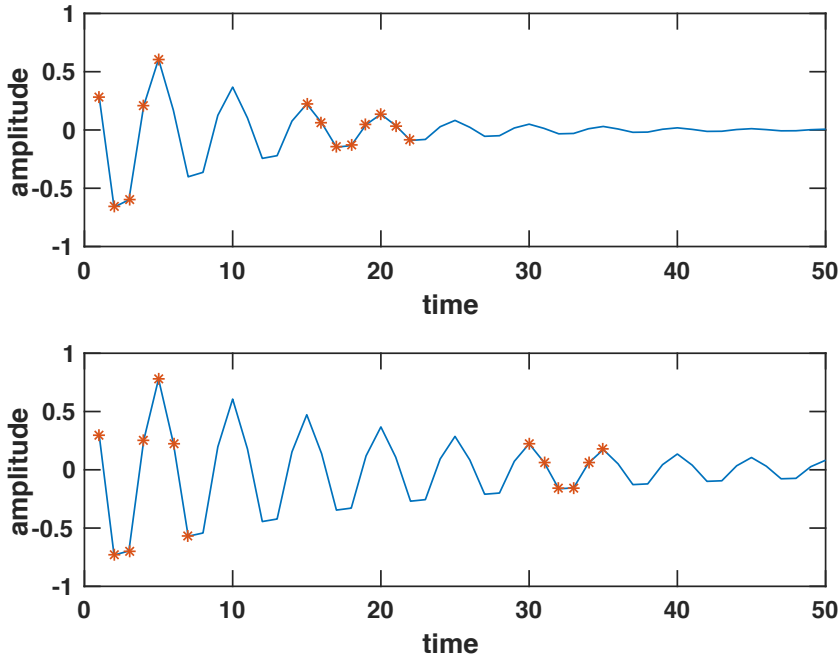

Figure 1: The resulting sampling scheme for two different values of β plotted against the real part of the signal. The upper most figure details the sampling scheme for $\beta = \frac{1}{10}$ and the bottom figure the sampling scheme for $\beta = \frac{1}{20}$.

solving (4) for two different levels of decay for a single damped sinusoid such that $\beta = 1/10$ for the top figure, and $\beta = 1/20$ for the bottom figure, but otherwise identical signal parameters. In both cases, $\gamma = 13$ so that $M = 13$ sample points, out of $N = 50$ possible candidates, are selected. Also, $\psi_p = 1$, $p = 1, \ldots, 4$, i.e., all signal parameters are considered in the minimization. As can be seen, the placing of the samples are determined by the damping parameter. As may be expected, for both values of β, some samples are placed in the beginning of the signal, where the SNR is at its maximum. To allow for an accurate estimation of

[1] The assumption of white circularly symmetric Gaussian noise is appropriate for the considered application to spectroscopy, where the noise is primarily resulting from thermal (Johnson) noise [32]. It is worth stressing that this is not a limitation for the sampling method as such, and the noise model may be selected to be non-Gaussian if deemed appropriate for the considered application; such a change would only affect the form of the used FIM.

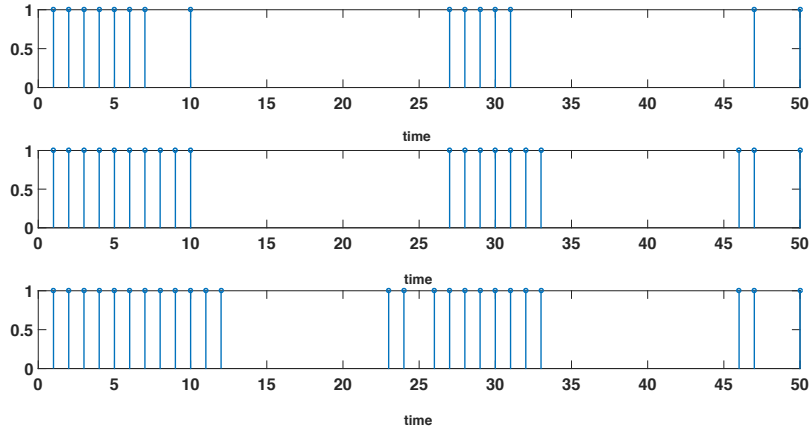

Figure 2: The resulting sample scheming for three different settings of γ, namely $\gamma = 15$, $\gamma = 20$, and $\gamma = 25$, where the signal contains two linear chirps.

the damping constant, one can also note that a further set of samples are selected later in the signal, with the more strongly decaying signal selecting them earlier than the less damped version, agreeing with the intuition that the more rapidly decaying signal contains less information at later sampling times.

As a further example, we next consider an example showing the resulting sample scheme for a signal containing two linear chirp components on the form

$$y(t_n) = \sum_{k=1}^{2} \alpha_k \exp\left(2i\pi\left(f_k^0 + f_k^1 t_n\right) t_n + i\varphi_k\right) + \varepsilon(t_n),$$

where f_k^0 and f_k^1 denote the frequency starting point and the slope of the chirp component k, respectively. Figure 2 shows the three sampling schemes yielded by the proposed method for three different setting on γ, namely $\gamma = 15$, $\gamma = 20$, and $\gamma = 25$. The here used parameters had the values $\alpha_1 = \alpha_2 = 5$, $f_1^0 = 0.1$, $f_2^0 = 0.5$, $f_1^1 = 0.01$, $f_2^1 = -0.003$, and the phases were set to $\varphi_1 = \pi/2$, and $\varphi_2 = \pi/3$. Due to the linear drift in frequency, it is reasonable to assume that the resulting sample scheme should have at least two clusters; one in the beginning of the signal, and one at the end of the signal. Looking at the sampling schemes in Figure 2 supports this intuition; three clusters are present for all three settings

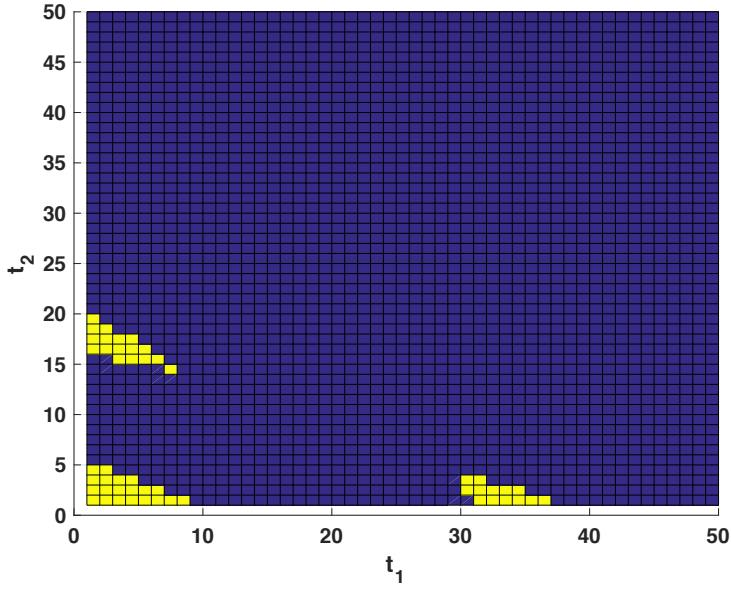

Figure 3: The resulting sampling scheme consisting of 50 selected samples for a signal consisting of a 2-D damped sinusoid as found when solving (4) with all $\psi_p = 1$.

of γ. When γ increases the two first clusters get bigger, whereas the last cluster remains more or less unchanged.

3.2 Illustration in 2-D

As further illustration of the impact of the choice of weight parameters ψ_p, consider the 2-D case with one damped sinusoid, i.e.,

$$y(t_1, t_2) = \alpha e^{2i\pi(f_1 t_1 + f_2 t_2) - (\beta_1 t_1 + \beta_2 t_2) + i\varphi} + \varepsilon(t_1, t_2), \tag{8}$$

with $\alpha = 1$, $f_1 = 0.2$, $f_2 = 0.5$, $\beta_1 = 1/20$, $\beta_2 = 1/10$, $\varphi = 1/2$, and noise variance $\sigma^2 = 0.1$. Figure 3 presents the sampling scheme found by solving (4) with $\gamma = 50$, i.e., 50 sampling points are chosen, for the case when $\psi_p = 1$ for all parameters.

As can be seen, the optimal sampling pattern here consists of three clusters of selected sampling points; one close to the origin and two close to the two time axes. Note that this is analogous to the 1-D case as the sampling cluster close to the first time axis is located further from the origin due to the decay in the first dimension being slower.

In contrast, Figure 4 displays the corresponding scheme found when solving (4), again with $\gamma = 50$, but only giving weight to the frequency and damping parameters, i.e., the ψ_p corresponding to the amplitude and phase parameters are set to zero. As can be seen, assigning the amplitude and phase parameters zero weight has the effect of shifting sampling points away from the origin to the clusters close to the t_1 and t_2 axes, in order to put more emphasis on the frequency and damping parameters. Indeed, the sum of the CRLBs for the parameters, as given by the sampling scheme in Figure 3, is $2.31 \cdot 10^{-2}$, whereas it is $3.61 \cdot 10^{-2}$ for the sampling scheme in Figure 4. However, if one considers the sum of the CRLBs for the frequency and damping parameters, these are $6.53 \cdot 10^{-4}$ and $4.42 \cdot 10^{-4}$ for Figures 3 and 4, respectively.

3.3 Simulations in 1-D

3.3.1 Optimization vs simulation

In Figure 5, we motivate that solving (6) is indeed a reasonable approach to determine optimal sampling patterns. The figure shows the obtained sum of the CRLBs for the parameters, i.e., $\operatorname{tr}\left(\mathcal{I}(\hat{\mathbf{w}}; \boldsymbol{\vartheta})^{-1}\right)$, where the sampling pattern is obtained by solving (6) for the case of $K = 1$ using the model (7), for a singleton set Θ. This is done for varying values of γ such that the number of samples used vary between $M = 5$ and $M = 25$. As a comparison, for each sample size M, we carry out 10^6 Monte Carlo simulations, in which we randomly decide on which M sampling points to use. We then compute which one of these 10^6 sampling patterns that results in the lowest sum of CRLBs. This will generally result in a sampling scheme significantly preferable to any fixed sampling scheme, such as Poisson gap sampling. As can be seen from the figure, the randomized approach achieves better results for small sample sizes, this as the simulations then become an exhaustive search, i.e., the simulations will with high likelihood find the exact solution to (1). However, as the sample size increases, so does the number of possible sampling patterns, which is $\binom{N}{M}$. As can be seen from the figure, the sampling scheme determined by (6) is then able to achieve an optimal perform-

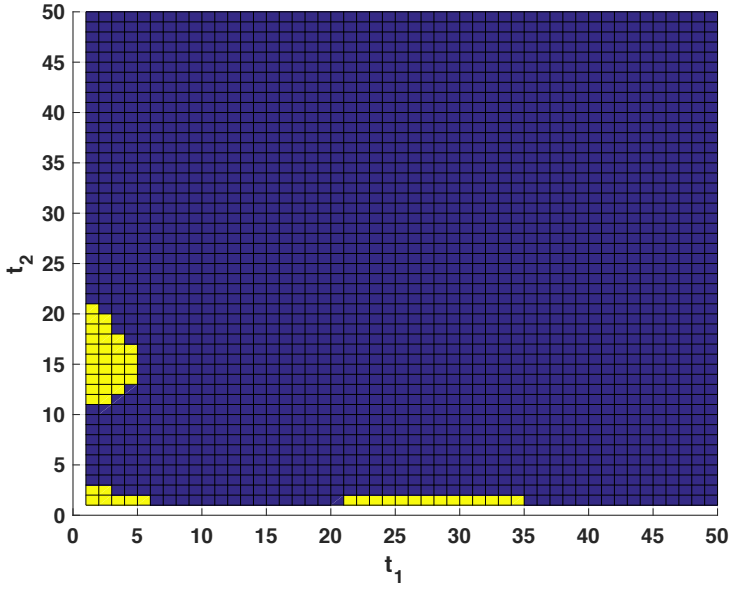

Figure 4: The resulting sampling scheme consisting of 50 selected samples for a signal consisting of a 2-D damped sinusoid as found when solving (4) with all $\psi_p = 1$ except for the amplitude and phase parameters, for which $\psi_p = 0$.

ance as the sample size increases.

3.3.2 Weighting

In Figures 6 and 7, we proceed to examine the effect of using different weightings for the signal parameters when solving (4). This is done for a signal consisting of two damped sinusoids with parameters $(\alpha_1, f_1, \beta_1, \varphi_1) = (1, 0.2, 1/12, 0.5)$ and $(\alpha_2, f_2, \beta_2, \varphi_2) = (1, 0.65, 1/20, \pi/5)$. The noise variance was $\sigma^2 = 0.01$ and $N = 50$. Assuming that we are interested only in the frequencies f_1, f_2, and the damping factors β_1, β_2, but not in the amplitudes or the phases, the weight parameters ψ_p are set to one for the frequency and damping parameters, whereas they are set to zero for the amplitudes and phases. Thus, the sought sampling pattern will be designed to increase the accuracy for the frequency and damping parameters at the expense of the amplitude and phase parameters.

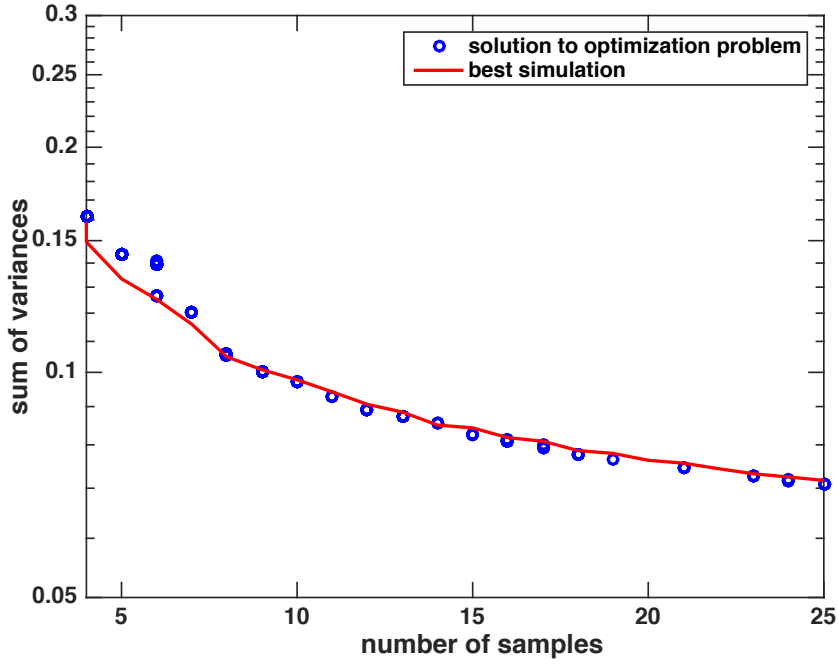

Figure 5: Sum of CRLBs for the parameters, i.e., tr $\left(\mathcal{I}(\hat{\mathbf{w}};\boldsymbol{\vartheta})^{-1}\right)$, for the sampling patterns given by the optimization problem and the best simulation, respectively, for different number of sampling points.

The resulting root CRLB, as a function of the number of samples used, for the frequencies f_1 and f_2 and the dampings β_1 and β_2 are shown in Figures 6 and 7, respectively. The root CRLB for the frequencies f_1 and f_2 is here defined as the root of the sum the individual CRLBs, and correspondingly for the dampings, β_1 and β_2. For comparison, the figures also present the root CRLBs corresponding to the optimal sampling patterns obtained for the case when no weighting is applied, i.e., $\psi_p = 1$ for all signal parameters. As can be seen, the weighting scheme results in sampling patterns that decrease the CRLB for the parameters of interest, in this case the frequencies and dampings. Also plotted is the obtained root mean squared error (RMSE) for the frequency and damping parameters, respectively, obtained when estimating these parameters using non-linear least squares (NLS) applied to simulated signals. The NLS estimate is found by solving

$$\hat{\boldsymbol{\vartheta}} = \underset{\boldsymbol{\vartheta}}{\operatorname{argmin}} \quad \frac{1}{2}\left\|\mathbf{y} - g(\boldsymbol{\vartheta})\right\|_2^2, \tag{9}$$

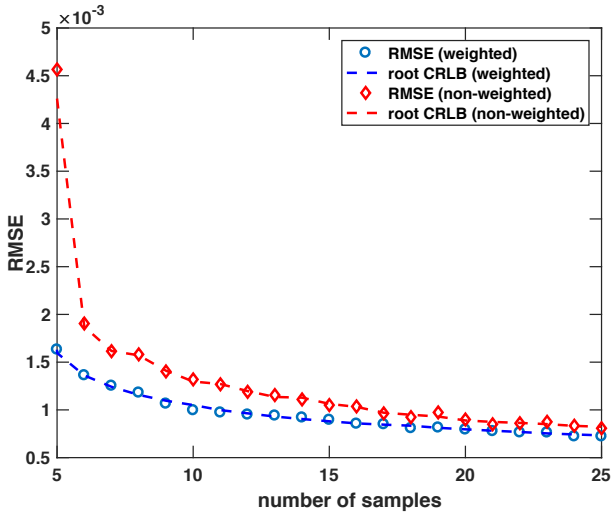

Figure 6: Obtained RMSE for the frequencies, when using the sampling patterns for the weighted and non-weighted cases, respectively.

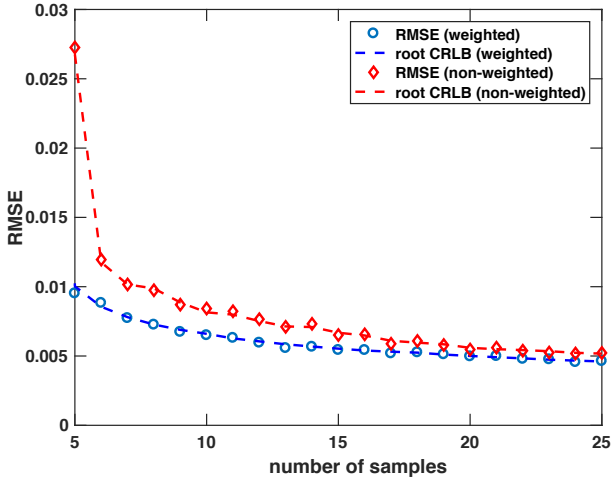

Figure 7: Obtained RMSE for the damping, when using the sampling patterns for the weighted and non-weighted cases, respectively.

where $\mathbf{y}$ is the data and $g(\boldsymbol{\vartheta})$ is the (non-linear) data model with parameter $\boldsymbol{\vartheta}$. In

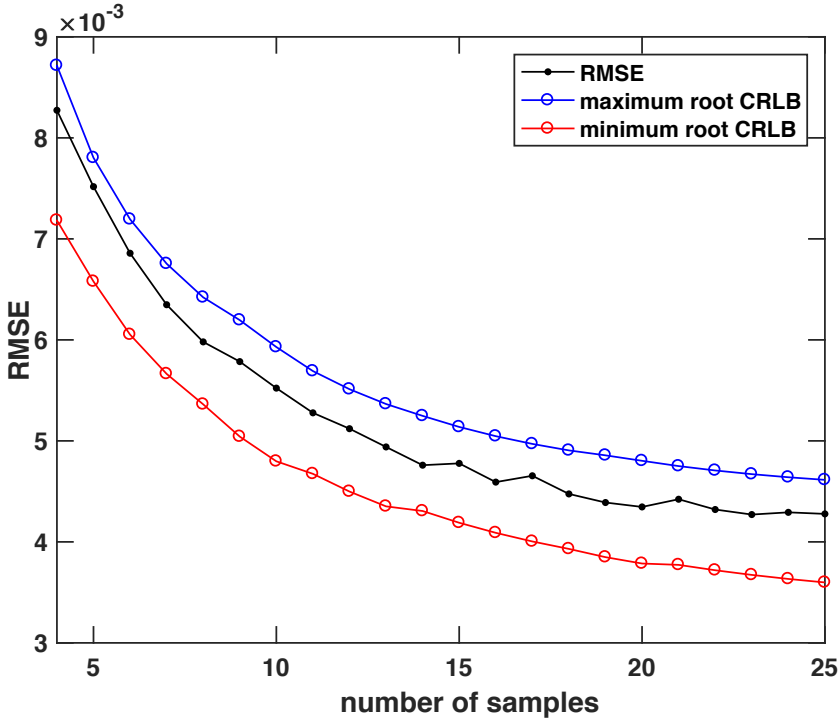

Figure 8: Obtained RMSE for the frequency f, when estimating $\boldsymbol{\vartheta}$ for the sampling pattern obtained for a grid of damping parameters β.

this paper, a minimum of (9) is found by evaluating the cost function over a grid of parameter values $\boldsymbol{\vartheta}$. The $\boldsymbol{\vartheta}$ that achieves the lowest value of (9) then becomes the resulting estimate. The RMSE is here defined as the root of the sum of the individual MSEs for the frequencies and dampings, respectively. As can be seen, the RMSE coincides with the root CRLB, implying that the bound is tight.

3.3.3 Gridding

Figures 8 and 9 show the effect of finding an optimal sampling pattern for a set of parameters $\boldsymbol{\vartheta} \in \Theta$ when solving (6). The results are obtained for a single decaying sinusoid. Here, we let $\Theta = \{\boldsymbol{\vartheta}_\ell\}_{\ell=1}^{L}$ express uncertainty in only the damping parameter β by fixing α, f, and φ, and letting Θ be a gridding over the damping parameter β, such that the parameter vectors constituting Θ are

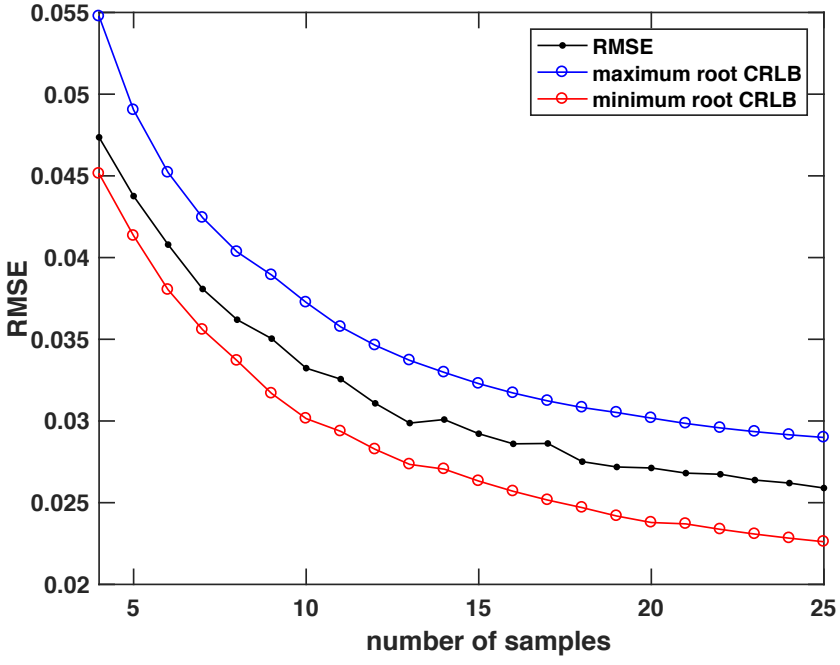

Figure 9: Obtained RMSE for the damping β, when estimating $\boldsymbol{\vartheta}$ for the sampling pattern obtained for a grid of damping parameters β.

$\boldsymbol{\vartheta}_\ell = (\alpha, f, \beta_\ell, \varphi)^T$, where $\beta_\ell = \beta_{\text{lower}} + \frac{\ell-1}{L}\Delta_\beta$ with Δ_β denoting the grid spacing, in effect letting β reside in the uncertainty interval

$$\mathcal{J}_\beta = \left[\beta_{\text{lower}}, \beta_{\text{lower}} + \frac{L-1}{L}\Delta_\beta\right].$$

The parameters used are $\alpha = 1$, $\varphi = 0.5$, $\sigma^2 = 0.1$, $\beta_{\text{lower}} = 0.1$, $\Delta_\beta = 0.022$, and $L = 10$. Using this, we solve (6) to get optimal sampling patterns as the number of samples grows. To evaluate the performance of the obtained sampling schemes, we then randomly sample the parameter vectors $\boldsymbol{\vartheta}$ where β is sampled uniformly on $\mathcal{J}_\beta$, i.e., on the interval covered by the grid Θ, but not on the grid points β_ℓ, $\ell = 0, 1, \ldots, L-1$. We then estimate $\boldsymbol{\vartheta}$ using NLS and compute the RMSE for the parameters $\boldsymbol{\vartheta}$. The figures show the obtained RMSE using 5000 Monte Carlo simulations for the frequency f and the damping β, respectively. Also presented are the best and worst case root CRLBs found on the grid Θ for

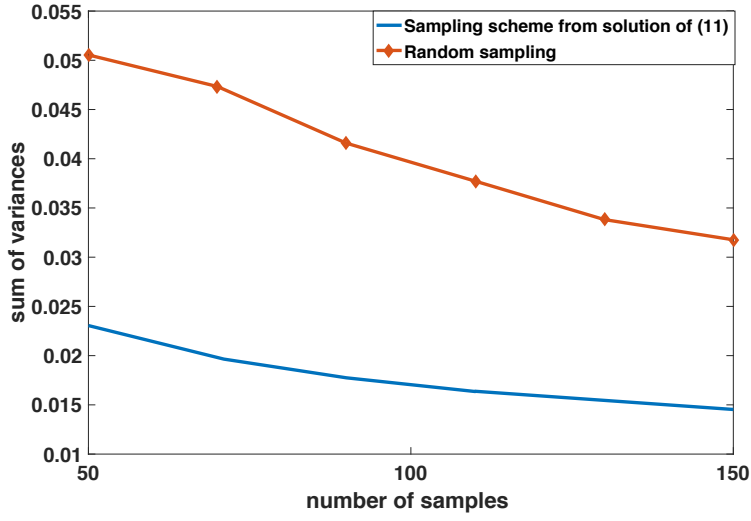

Figure 10: The sum of variances of the parameters of interest as a function of the number of selected samples.

each parameter. The obtained RMSE lies between the lowest and highest on-grid root CRLB for both parameters and for all considered sample sizes, suggesting that (6) indeed yields sampling schemes with a guaranteed worst case performance, as well as a lower limit on the possible RMSE.

3.4 Simulations in 2-D

3.4.1 Optimization vs simulation

As was seen in the 1-D setting, the optimization scheme was able to outperform the method of randomly selecting sampling points and then choosing the scheme minimizing the sum of the CRLB. In 2-D, this becomes even more apparent as the number of potential sampling points increase rapidly with increasing dimension. An illustration of this is shown in Figure 10, showing the sum of the CRLBs obtained when solving (6) for varying numbers of desired sampling points. The signal considered is the 2-D damped sinusoid in (8) with parameters $\alpha = 1, f_1 = 0.2, f_2 = 0.5, \beta_1 = 1/20, \beta_2 = 1/10, \varphi = 1/2$, and $\sigma^2 = 0.1$. We here let $\psi_p = 1$ for all signal parameters, and consider a sampling space of 50×50 potential

sampling times. Also presented is the sum of the CRLBs for the best (defined as the one with smallest sum of CRLBs) among 10^7 sampling scheme obtained by randomly choosing sampling points. As can be seen from the figure, the proposed method outperforms the random sampling for all numbers of selected samples. It is worth noting that the computational time to evaluate the 10^7 sampling schemes was three times longer than solving the proposed problem using a off-the-shelf convex solver [33].

3.4.2 Weighting

We here consider the case of a signal consisting of two 2-D damped sinusoid, i.e.,

$$y(t_1, t_2) = \sum_{k=1}^{K} \alpha_k e^{i\varphi_k} \Pi_{d=1}^{2} e^{2i\pi f_{k,d} t_d - \beta_{k,d} t_d} + \varepsilon(t_1, t_2),$$

for $K = 2$. Let the parameters be $(f_{1,1}, f_{2,1}) = (0.1, 0.2)$ and $(\beta_{1,1}, \beta_{2,1}) = (0.1, 0.1)$ for the first dimension, $(f_{1,1}, f_{2,1}) = (0.1, 0.2)$ and $(\beta_{1,1}, \beta_{2,1}) = (0.1, 0.1)$ for the second dimension, and let $\alpha_1 = 1$, $\alpha_2 = 1.3$, $\varphi_1 = \frac{\pi}{3}$, $\varphi_2 = \frac{\pi}{3}$, and $\sigma^2 = 0.01$. We then determine optimal sampling schemes by solving (4) for varying numbers of sampling points. This is done for both the equally weighted case, i.e., with $\psi_p = 1$ for all signal parameters, as well as for the case when only the frequency and damping parameters are given weight, i.e., with $\psi_p = 0$ for the amplitude and phase parameters. The results are shown in Figures 11-14. In Figure 11, the root of the sum of the CRLBs for the frequencies in the first dimension, i.e., $f_{1,1}$ and $f_{2,1}$, is shown. Similarly, Figure 12 corresponds to the frequencies in the second dimension, while Figures 13 and 14 correspond to the damping parameters in the first and second dimension, respectively. Also presented is the corresponding RMSE obtained when estimating the parameters using NLS. As can be seen, the obtained RMSEs coincides with the CRLBs for both the weighted and non-weighted case, implying that the bound is tight. Note also that the schemes corresponding to assigning no weight to the amplitude and phase parameters all result in a lower sum of CRLB for the frequency and damping parameters than the non-weighted schemes. This comes at the price of a larger sum of CRLB for the amplitudes α_1 and α_2, which is illustrated in Figure 15. As can be seen in the figure, the non-weighted sampling scheme here leads to more accurate estimates of the amplitudes.

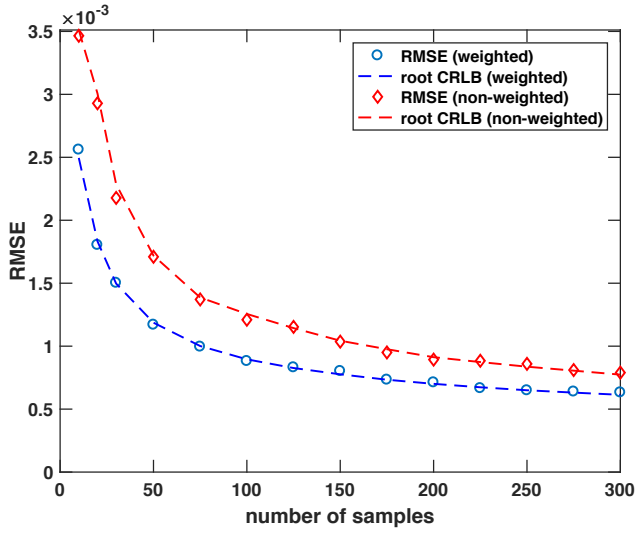

Figure 11: Obtained RMSE for the frequencies in the first dimension, when using the sampling patterns for the weighted and non-weighted cases, respectively.

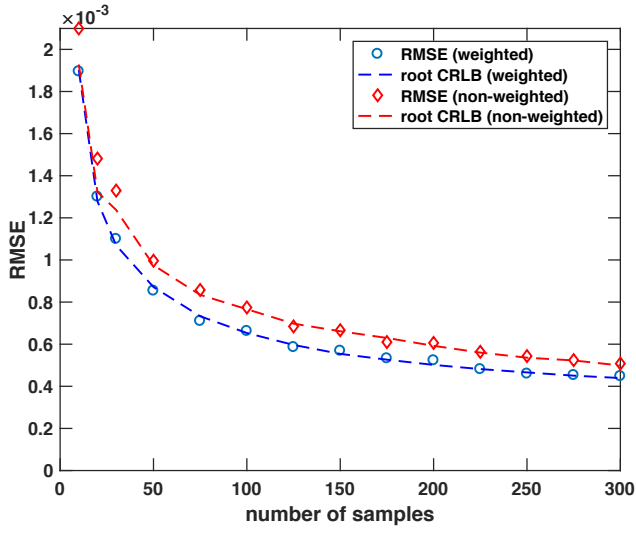

Figure 12: Obtained RMSE for the frequencies in the second dimension, when using the sampling patterns for the weighted and non-weighted cases, respectively.

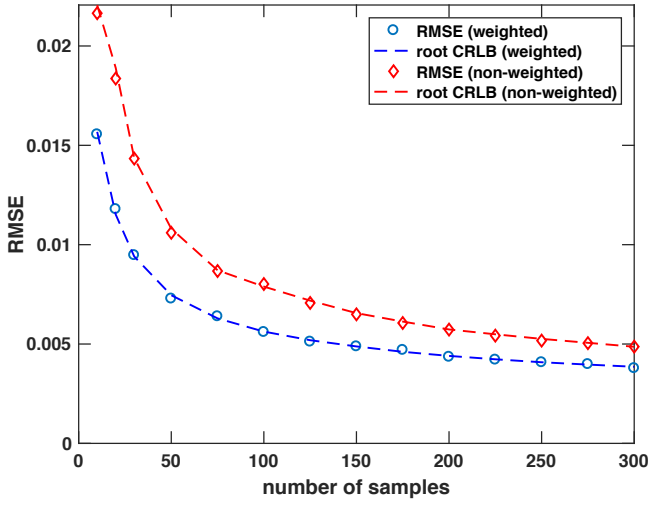

Figure 13: Obtained RMSE for the dampings in the first dimension, when using the sampling patterns for the weighted and non-weighted cases, respectively.

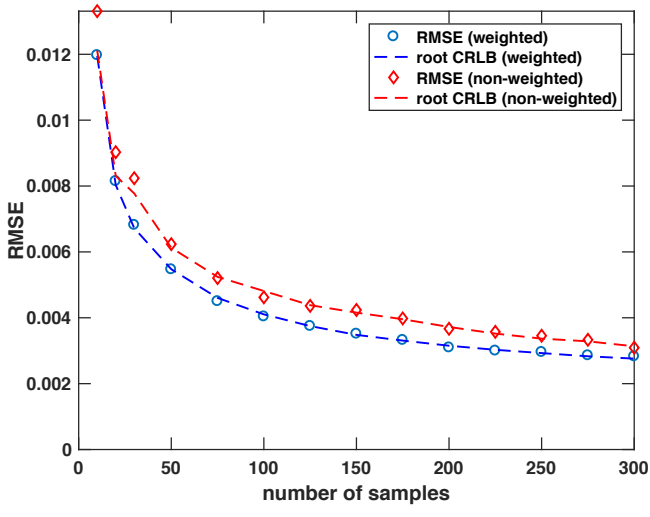

Figure 14: Obtained RMSE for the dampings in the second dimension, when using the sampling patterns for the weighted and non-weighted cases, respectively.

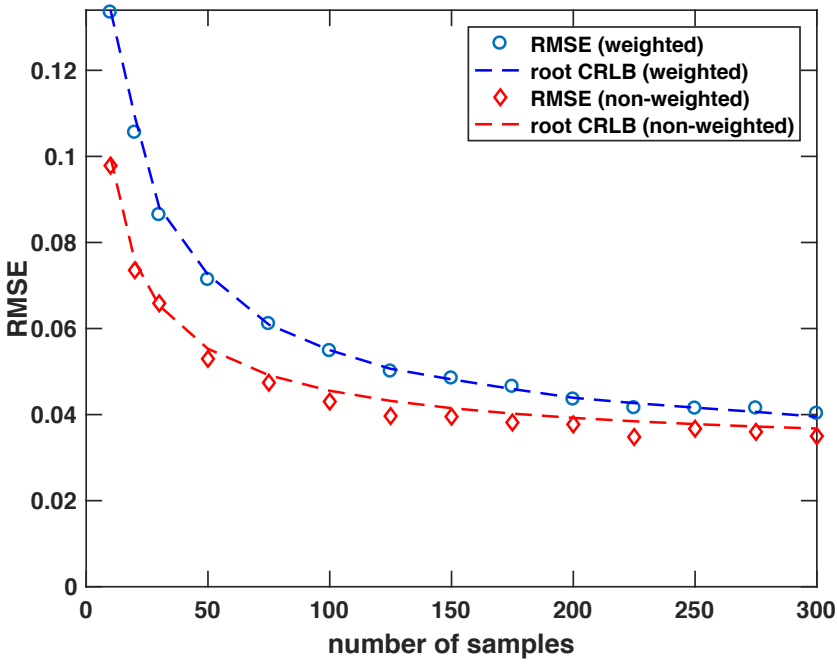

Figure 15: Obtained RMSE for the amplitudes, when using the sampling patterns for the weighted and non-weighted cases, respectively.

3.5 Real NMR signal

In order to illustrate the proposed framework's applicability to real, measured data, we first consider a 2-D NMR signal obtained from a ^{15}N-HSQC experiment made on a Histidine sample. The signal is sampled on a uniform 25×25 grid. In this form of experiments, the spectrometer measures the signal resulting from repeatedly pulsing the studied substance, each pulse resulting in a so-called free-induction decay (FID). This measurement is made in the direct dimension, i.e., along time. By modifying the pulse sequence, for instance changing the timing between pulses, one obtain measurements also in a second (indirect) dimension. By further modifications to the pulse sequence, one may similarly obtain higher dimensional data sets, where each dimension corresponds to the range over each of the considered pulse settings [34]. Figure 16 shows the 2-D periodogram of the signal using all available samples. As can be seen, the signal consists of a num-

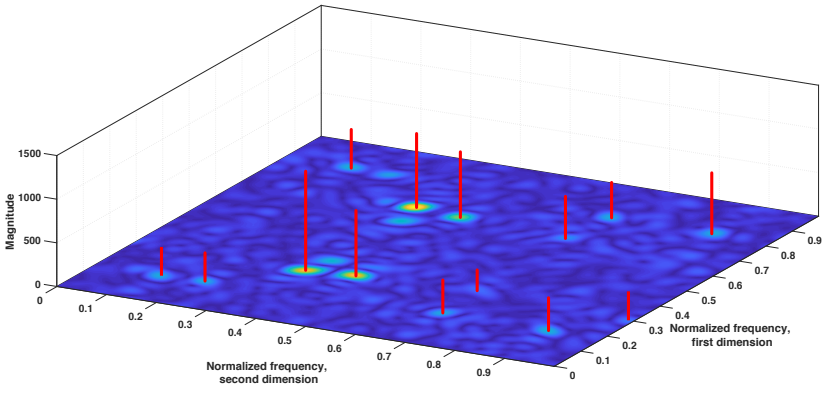

Figure 16: Estimates obtained using SEMA when applied to a real NMR signal using all available samples, superimposed on the corresponding periodogram estimate.

ber of components of varying powers, all which may be fairly well modeled as damped complex sinusoids. Superimposed on the periodogram are estimates obtained using the SEMA method presented in [35], also using all available samples. This method is a sparse reconstruction algorithm requiring no prior knowledge of the number of components constituting the signal. Considering only the four most prominent components estimated by SEMA, we then compute a sampling scheme consisting of 225 samples in total, i.e., a mere 36% of the number of available samples, using the formulation in (4), and estimate the number of signal components, as well as the parameters of the components, using SEMA. The result is presented in Figure 17, displaying the periodogram using all samples superimposed with the obtained SEMA estimates. Comparing Figures 16 and 17, it may be noted that the considerably reduced sample size does not cause any significant degradation in estimation performance.

Proceeding, we examine the generation of a 3-D sampling scheme for a laser spectroscopy measurement containing 4 damped sinusoids. The measured data set contains $40 \times 40 \times 20$ samples, out of which we select a total of 80 samples using (4). Figure 18 shows the resulting sampling scheme. The figure again illustrates, reminiscent to Figures 1 and 3, that samples are selected where the local SNR is highest, to allow for accurate frequency estimation, as well as a smaller

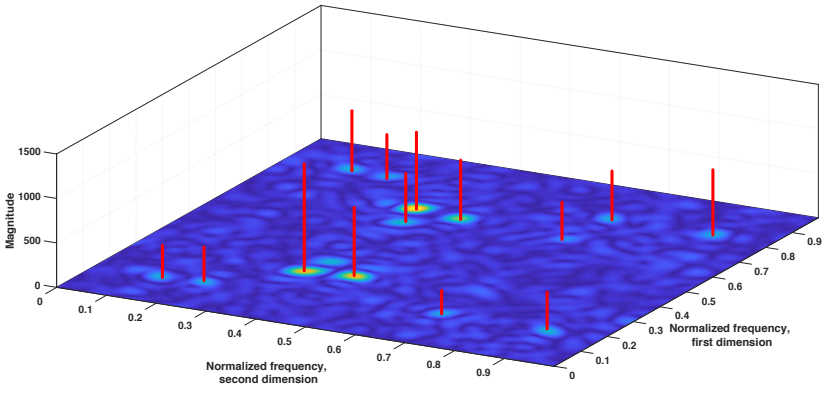

Figure 17: Estimates obtained using SEMA when applied to a real NMR signal 36% of the available samples as selected by (4), superimposed on the periodogram estimate using all available samples.

number of additional samples selected later in the sequences to better capture the damping behavior.

4 Conclusion

In this work, we have proposed a convex optimization problem for finding suitable sampling schemes for multidimensional data models. The optimization problem is formed as to, from a set of available samples, determine a subset of given cardinality such that linear combinations of the variances of the signal parameters of interest are minimized. Due to the structure of the optimization problem, it is easy to add additional constraints, e.g., adding performance bounds on selected parameters, or putting more emphasize on a subset of the parameters, and to model for the uncertainty in *a priori* assumptions of the parameter values. In the numerical section, we show that solving the proposed optimization problem is a more efficient approach than randomly selecting the sampling points, especially in the multi-dimensional setting. Further, we show that using the sampling schemes found by solving the proposed optimization problem, will provide a lower Cramér-Rao lower bound than that found from using ordinary uniform sampling. By using an efficient parameter estimator on the signal

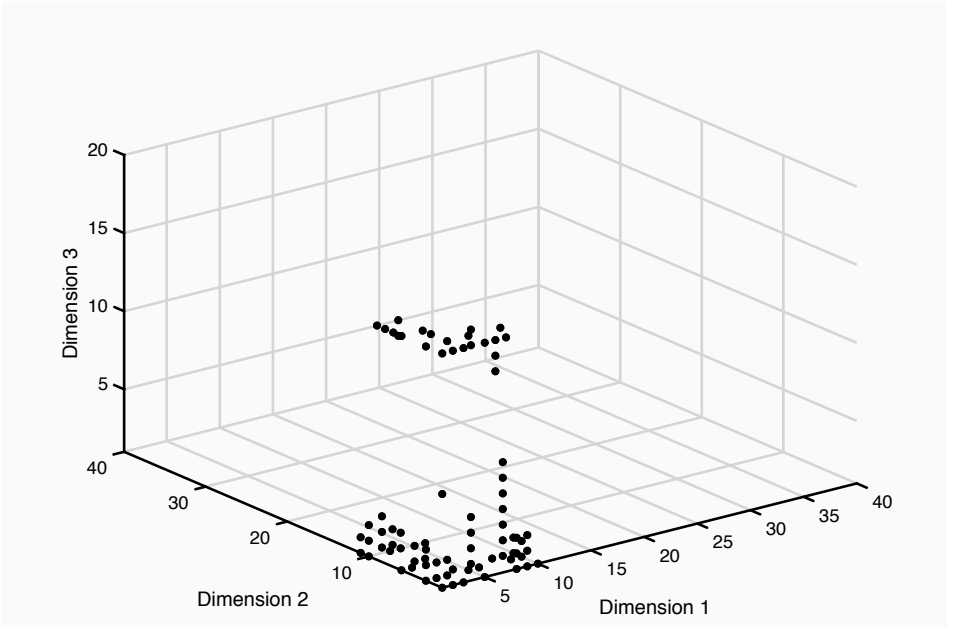

Figure 18: The resulting sampling scheme for a 3-D laser spectroscopic signal containing 4 damped sinusoids.

sampled according to the found sampling scheme, we show that these Cramér-Rao lower bounds are, in fact, tight.

5 Acknowledgement

The authors wish to express their gratitude to Prof. Mikael Akke, Prof. Tönu Pullerits, Dr Göran Carlström, and Dr Khadga Jung Karki, all at Lund University, for providing us with the here examined 2-D NMR and 3-D laser spectroscopy measurements.

References

[1] S. Liu, M. Fardad, E. Masazade, and P. K. Varshney, "Optimal Periodic Sensor Scheduling in Networks of Dynamical Systems," *IEEE Trans. Signal Process.*, vol. 62, no. 12, pp. 3055–3068, December 2014.

[2] H. Jamali-Rad, A. Simonetto, X. Ma, and G. Leus, "Distributed Sparsity-Aware Sensor Selection," *IEEE Trans. Signal Process.*, vol. 63, no. 22, pp. 5951–5964, 11 2015.

[3] S. P. Chepuri, G. Leus, and A. J. van der Veen, "Sparsity-Exploiting Anchor Placement for Localization in Sensor Networks," in *21st European Signal Processing Conference*, 9-13 September 2013, pp. 1–5.

[4] S. Ravishankar and Y. Bresler, "Adaptive Sampling Design for Compressed Sensing MRI," in *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Boston, Massachusetts, 30 Aug.-3 Sept 2011, pp. 3751–3755.

[5] F. Gama, A. G. Marques, G. Mateos, and A. Ribeiro, "Rethinking Sketching as Sampling: Linear Transformation of Graph Signals," in *50th Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, November 2016.

[6] A. Anis, A. Gadde, and A. Ortega, "Efficient Sampling Set Selection for Bandlimited Graph Signals Using Graph Spectral Proxies," *IEEE Trans. Signal Process.*, vol. 64, no. 14, pp. 3775–3789, July 2016.

[7] P. Schmieder, A. S. Stern, G. Wagner, and J. C. Hoch, "Application of nonlinear sampling scheme to COSY-type spectra," *Journal of Biomolecular NMR*, vol. 3, pp. 569–576, 1993.

[8] K. Kazimierczuk, A Zawadzka-Kazimierczuk, and W. Koźmiński, "Non-uniform frequency domain for uniform exploitation of non-uniform sampling," *J. Magn. Reson.*, vol. 205, pp. 286–292, 2010.

[9] S. G. Hyberts, K. Takeuchi, and G. Wagner, "Poisson-Gap Sampling and Forward Maximum Entropy Reconstruction for Enhancing the Resolution and Sensitivity of Protein NMR Data," *J Am Chem Soc.*, vol. 132, pp. 2145–2147, 2010.

[10] P. J. Sidebottom, "A new approach to the optimisation of non-uniform sampling schedules for use in the rapid acquisition of 2D NMR spectra of small molecules," *Magn Reson Chem*, vol. 54, no. 8, pp. 689–694, August 2016.

[11] K. Kazimierczuk and V. Y. Orekhov, "Accelerated NMR Spectroscopy by Using Compressed Sensing," *Angewandte Chemie International Edition*, vol. 50, no. 24, June 2011.

[12] K. Kazimierczuk and V. Y. Orekhov, "A comparison of convex and non-convex compressed sensing applied to multidimensional NMR," *J. Magn. Reson.*, vol. 223, pp. 1–10, 2012.

[13] M. Billeter, "Non-uniform sampling in biomolecular nmr," *Journal of Biomolecular NMR*, vol. 68, no. 2, pp. 65–55, 2017.

[14] S. G. Hyberts, H. Arthanari, S. A. Robson, and G. Wagner, "Perspectives in magnetic resonance: NMR in the post-FFT era," *J. of Magn. Reson.*, vol. 241, pp. 60–73, 2014.

[15] P. C. Aoto, R. B. Fenwick, G. J. A. Kroon, and P. E. Wright, "Accurate Scoring of Non-uniform Sampling Schemes for Quantitative NMR," *Journal of Magnetic Resonance*, vol. 246, pp. 31–35, Sept 2014.

[16] A. Özcelikkale, H. M. Ozaktas, and E. Arikan, "Signal Recovery with Cost-Constrained Measurements," *IEEE Trans. Signal Process.*, vol. 58, no. 7, pp. 3607–3617, July 2010.

[17] Y. Yilmaz, S. Li, and X. Wang, "Sequential Joint Detection and Estimation: Optimum Tests and Applications," *IEEE Trans. Signal Process.*, vol. 64, no. 20, pp. 5311, Oct 2016.

[18] V. Kekatos, G. B. Giannakis, and B. Wollenberg, "Optimal Placement of Phasor Measurement Units via Convex Relaxation," *IEEE Trans. on Power Systems*, vol. 27, no. 3, pp. 1521–1530, Aug 2012.

[19] T. C-Gulcu and H. M. Ozaktas, "Choice of Sampling Interval and Extent for Finite-Energy Fields," *IEEE Trans. Signal Process.*, vol. 65, no. 7, pp. 1741–1751, April 2017.

[20] H. Zhang, J. M. F. Moura, and B. K. Krogh, "Dynamic Field Estimation Using Wireless Sensor Networks: Tradeoffs Between Estimation Error and Communication Cost," *IEEE Trans. Signal Process.*, vol. 57, no. 6, pp. 2383–2395, June 2009.

[21] S. Liu, E. Masazade, and P. K. Varshney, "Temporally Staggered Sensing for Field Estimation with Quantized Data in Wireless Sensor Networks," in *IEEE Statistical Signal Processing Workshop (SSP)*, Ann Arbor, MI, USA, August 2012.

[22] S. Joshi and S. Boyd, "Sensor Selection via Convex Optimization," *IEEE Trans. Signal Process.*, vol. 57, no. 2, pp. 451–462, 2 2009.

[23] F. Pukelsheim, *Optimal design of experiments*, Wiley series in probability and mathematical statistics. Wiley, New York, 1993.

[24] S. P. Chepuri, *Sparse Sensing for Statistical Inference - Theory, Algorithms, and Applications*, Ph.D. thesis, Delft University of Technology, 2015.

[25] S. P. Chepuri and G. Leus, "Sparsity-Promoting Sensor Selection for Non-Linear Measurement Models," *IEEE Trans. Signal Process.*, vol. 63, no. 3, pp. 684–698, 2 2015.

[26] S. Liu, S. P. Chepuri, M. Fardad, E. Masazade, and G. Leus P. K. Varshney, "Sensor Selection for Estimation with Correlated Measurement Noise," *IEEE Trans. Signal Process.*, vol. 64, no. 13, pp. 3509–3522, July 2016.

[27] S. P. Chepuri and G. Leus, "Continuous Sensor Placement," *IEEE Signal Process. L*, vol. 22, no. 5, pp. 544–548, May 2015.

[28] J. Swärd, F. Elvander, and A. Jakobsson, "Designing Optimal Sampling Schemes," in *25th European Signal Processing Conference*, Aug 28 - Sep 2 2017.

[29] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004.

[30] J. F. Sturm, "Using SeDuMi 1.02, a Matlab toolbox for optimization over symmetric cones," *Optimization Methods and Software*, vol. 11-12, pp. 625–653, August 1999.

[31] R. H. Tutuncu, K. C. Toh, and M. J. Todd, "Solving semidefinite-quadratic-linear programs using SDPT3," *Mathematical Programming Ser. B*, vol. 95, pp. 189–217, 2003.

[32] J. B. Johnson, "Thermal Agitation of Electricity in Conductors," *Phy. Rev.*, vol. 32, pp. 97–109, July 1928.

[33] Inc. CVX Research, "CVX: Matlab Software for Disciplined Convex Programming, version 2.0 beta," `http://cvxr.com/cvx`, Sept. 2012.

[34] J. Keeler, *Understanding NMR Spectroscopy*, Chichester, U.K. : John Wiley and Sons, 2010.

[35] J. Swärd, S. I. Adalbjörnsson, and A. Jakobsson, "High Resolution Sparse Estimation of Exponentially Decaying N-dimensional Signals," *Elsevier Signal Processing*, vol. 128, pp. 309–317, Nov 2016.