

JAKOB MÖDERL

DETECTION AND ESTIMATION OF DISPERSIVE TARGET SIGNALS





Dipl.-Ing. Jakob Möderl, BSc

Detection and Estimation of Dispersive Target Signals

DOCTORAL THESIS

to achieve the university degree of
Doktor der technischen Wissenschaften
submitted to

Graz University of Technology

Supervisors

Univ.-Prof. Dipl.-Ing. Dr. Klaus Witrals
Dipl.-Ing. Dr.techn. Erik Leitinger

Institute of Communication Networks and Satellite Communications

Graz, September 2024

Acknowledgment

Although only one name appears on the cover, the completion of this thesis would not have been possible without the direct and indirect input of numerous individuals. I would like to acknowledge the significant influence that my academic tutors and supervisors have had on my academic career, without whom I would not have been in a position to complete this thesis. Particularly, I want to express my gratitude to my supervisor Klaus Witrals. I greatly appreciate your continued support. Even more so, I appreciate the autonomy afforded to me to pursue my own ideas. I also greatly appreciate my co-supervisor Erik Leitingner for taking an interest in my work. This thesis would have gone in quite a different (and presumably less interesting) direction without your invaluable input. Science and research thrives on curiosity combined with creative and critical thinking. These skills are best exercised in friendly discussions. Therefore, I want to thank all colleagues and friends at the IKS and SPSC who engaged with me in countless discussions over the years, whether on technical, philosophical or personal matters.

The second group of individuals who provided indirect but indispensable contributions to this thesis are those who encouraged me to embark on this endeavour and provided me with support throughout its duration. I am certain that I would never have considered pursuing a PhD without the encouragement and guidance of my parents, Susanne and Ulrich, who fostered my interest in knowledge and inquiry during my formative years. Furthermore, my parents and grandparents enabled me to attend university and provided me with financial support throughout my studies, for which I am forever grateful. Finally, but most importantly, I want to thank my long-time girlfriend Patricia. Whenever my thoughts get tangled erring through the abstract worlds of sparsity, spectral lines, probability distributions, and variational Bayesian inference, you always manage to untangle them by pulling me back into reality. Thank you!

Abstract

The use of ultra-wideband (UWB) radios enables the acquisition of high-accuracy time-of-flight information, a consequence of their high bandwidth. Thus, UWB radars are superior at resolving closely-spaced multipath components (MPCs) relative to narrowband radars, particularly in indoor environments. However, due to the high time resolution, it is necessary to consider large targets, e.g. people and cars, as extended objects that disperse the backscattered signal over time. To account for such dispersion, two hypotheses are investigated in this thesis.

The first hypothesis proposes to utilize the dispersiveness of indoor multipath environments to enhance the detection of occupants. This study specifically investigates using UWB radars to detect car occupants based on their respiratory motion. By employing a point-target assumption and pinhole channel model, we show that the respiratory motion and multipath channel response can be factorized. Two detection algorithms are derived from this factorized model: an estimator-correlator (EC) and a detector based on variational inference (VI). Both of these detectors are shown to outperform the baseline fast Fourier transform (FFT) approach, which assumes a line-of-sight-only model. Furthermore, a data-driven approach is also applied to this use case based on the widely-used ResNet architecture. A measurement campaign was conducted to obtain a dataset for training the network. The dataset comprises radar channel impulse responses (CIRs) of participants engaged in a range of motion activities within the vehicle. The data-driven detection algorithm outperforms the model-based approaches. However, since the dataset contains only data from adults, the ResNet’s generalization capabilities should be validated on cases not contained in the dataset, such as infants or small children.

The second hypothesis proposes to enhance the detection and estimation performance of sparsity-based line spectral estimation (LSE) for structured line spectra using group-sparsity. Structured line spectra, i.e. spectra comprised of clusters of related spectral lines, arise in several applications. E.g., the time-dispersed backscattered signal from extended objects is approximated by a structured line spectrum. However, the majority of LSE algorithms assume the spectral lines to be unstructured, i.e. they assume all spectral lines are uncorrelated. We derive two sparse Bayesian learning (SBL)-based LSE algorithms which utilize group-sparsity to enhance the estimation performance for structured line spectra. The fast variational block-SBL algorithm is based on the assumption that the group-size is known and the dictionary fixed. The more sophisticated group-sparse super-resolution algorithm incorporates the estimation of dictionary parameters (on a continuum) and the group sizes into the estimation procedure. Through simulations, we show that our group-sparse super-resolution algorithm increases the detection probability of the spectral lines related to extended objects in low-signal-to-noise ratio (SNR) scenarios. As a consequence, the parameters of the extended object, e.g. its center

of mass and extent, can be estimated more accurately. Furthermore, structured line spectra are also observed in other applications. We apply the fast variational block-SBL algorithm to multiple measurement vector (MMV) direction of arrival (DOA) estimation using data from an antenna array, and the group-sparse super-resolution algorithm to multi-pitch estimation of audio data. The algorithms presented in this thesis outperform the current state-of-the-art methods in both cases.

In summary, this thesis contributes to the detection and estimation of dispersive target signals. It does so both practically, by developing detection algorithms utilizing the dispersiveness of dense multipath channels to enhance the detection performance in the car occupancy detection use case, and theoretically, by developing LSE algorithms tuned to the detection and estimation of structured line spectra.

Preface: About Models

“All models are wrong, but some are useful.”

—George E. P. Box [1]

Mathematical models are central to Bayesian statistics, e.g. a likelihood function describing how observed data depend on unknown parameters of interest. Nevertheless, it is not possible for any model, whether statistical, of a radio system, or otherwise, to describe the intricate details of our macro-realm reality with absolute precision. In other words, every model is an approximation. If the underlying assumptions of a model are not fulfilled, a mathematician may conclude that the model is incorrect. Nevertheless, this may be of little consequence in light of the fact that all models are, to some extent, imperfect representations of reality.

The aforementioned quote succinctly expresses that the merit of a model should be evaluated based on its efficacy in fulfilling the intended purpose, rather than on its precision in representing reality. In other words, it expresses the engineering principle of *“it’s not stupid if it works”*. Moreover, according to the principle of Occam’s razor, the objective is to identify the most straightforward description that encompasses the essential elements of reality, providing sufficient detail for the specific task at hand. Following this line of reasoning, each model presented in this thesis has been designed with the practical purpose of developing estimation and/or detection algorithms. The presented models are not designed to be the most accurate or complex models that describe reality in the greatest detail possible. Rather, they are deliberately designed to have a mathematical structure that can be utilized in a straightforward and meaningful manner to facilitate the development of estimation and/or detection algorithms.

Contents

I	Introduction and Main Results	1
1	Introduction	3
1.1	Research Hypotheses	4
1.1.1	Car Occupancy Detection Using Ultra-Wideband Radar . . .	4
1.1.2	Estimation of Structured Line Spectra	4
1.2	Contribution and Outlook	5
2	Background	7
2.1	Toy Model of a Radio System	8
2.1.1	System Model	8
2.1.2	Channel Model	10
2.2	Line Spectral Estimation (LSE)	11
2.2.1	Sparse Bayesian Learning (SBL)	16
2.3	Variational Inference (VI)	18
2.3.1	The Variational Bayesian Approach	18
2.3.2	Mean-Field Approximation	19
2.3.3	Variational Expectation-Maximization (EM) Algorithm . . .	22
3	Car Occupancy Detection	25
3.1	Model-driven Car Occupancy Detection	26
3.1.1	Model of the Empty Car	26
3.1.2	Target Model	27
3.1.3	Energy Detector (ED)	29
3.1.4	FFT-Detector	30
3.1.5	Estimator-Correlator (EC)	31
3.1.6	Variational Inference (VI)	32
3.2	Data-driven Car Occupancy Detection	34
3.2.1	Dataset Description	34
3.2.2	ResNet-based Occupancy Detection	35
3.3	Results	38
3.4	Concluding Remarks	39
4	Estimation of Structured Line Spectra	41
4.1	Modeling Extended Objects Using Structured Line Spectra	42
4.2	Fast Inference With Known Group Sizes	44
4.2.1	Introduction and Related Methods	44
4.2.2	Fast Variational Block SBL	46

4.2.3	Results	48
4.2.4	Application: Multiple Measurement Vector (MMV)-based Direction-of-Arrival (DOA) Estimation	49
4.3	Super-Resolution Inference With Unknown Group Sizes	51
4.3.1	Variational-EM Approach	51
4.3.2	Application: Extended Object Detection	55
4.3.3	Application: Multi-Pitch Estimation	57
4.4	Concluding Remarks	60
5	Conclusion	61
II	Included Papers	71
A	Car Occupancy Detection Using Ultra-Wideband Radar	73
B	Variational Message Passing-based Respiratory Motion Estimation and Detection Using Radar Signals	83
C	UWBCarGraz Dataset for Car Occupancy Detection using Ultra- Wideband Radar	93
D	Fast Variational Block-Sparse Bayesian Learning	107
E	Variational Inference of Structured Line Spectra Exploiting Group- Sparsity	137

Notation and Common Functions

\mathbf{a}, \mathbf{A}	Bold letters \mathbf{a} or \mathbf{A} are used for vectors and matrices while non-bold letters a and A are used for scalars.
$[\mathbf{x}_l : l \in \mathcal{L}]$	The matrix formed by the columns \mathbf{x}_l for all $l \in \mathcal{L} \subset \mathbb{Z}$, ordered by increasing l .
$[\mathbf{a}]_i, [\mathbf{A}]_{i,j}$	The operator $[\mathbf{a}]_i$ returns the i th element of the vector \mathbf{a} . Similarly, the operator $[\mathbf{A}]_{i,j}$ returns the element in the i th column and j th row of the matrix \mathbf{A} .
$\ \mathbf{a}\ _p, \ \mathbf{a}\ $	$\ \mathbf{a}\ _p$ denotes the ℓ^p -norm of the vector \mathbf{a} . $\ \boldsymbol{\alpha}\ \triangleq \ \mathbf{a}\ _2$ denotes the ℓ^2 norm.
$ \mathcal{S} $	Cardinality of set \mathcal{S} .
$\mathcal{S} \setminus a$	Shorthand notation for $\mathcal{S} \setminus \{a\}$.
$e^{\mathbf{a}}$	The function $e^{\mathbf{a}}$, where $\mathbf{a} = [a_1 \ a_2 \ \dots \ a_N]^T$ is a vector, denotes the element-wise exponential, i.e. $e^{\mathbf{a}} = [e^{a_1} \ e^{a_2} \ \dots \ e^{a_N}]^T$.
$ \mathbf{A} , \text{tr}(\mathbf{A})$	Determinant and trace of matrix \mathbf{A} .
$\text{vec}(\mathbf{A})$	The operator $\text{vec}(\mathbf{A})$ stacks all the columns \mathbf{a}_i of the matrix $\mathbf{A} = [\mathbf{a}_1 \ \mathbf{a}_2 \ \dots \ \mathbf{a}_N]$ into a long vector, i.e. $\text{vec}(\mathbf{A}) = [\mathbf{a}_1^T \ \mathbf{a}_2^T \ \dots \ \mathbf{a}_N^T]^T$.
$\mathbf{A} \otimes \mathbf{B}$	Kronecker product between vectors/matrices \mathbf{A} and \mathbf{B} .
$\mathbf{I}, \mathbf{0}$	The symbols \mathbf{I} and $\mathbf{0}$ denote the identity matrix and a vector/matrix of all zeroes, respectively. The dimensions are denoted by a subscript in case they are not clear from context, i.e. \mathbf{I}_N for the $N \times N$ identity matrix.
$\text{Re}\{\cdot\}, \text{Im}\{\cdot\}$	Operators returning the real and imaginary parts of the input.
$\mathcal{F}\{x(t)\}$	Fourier transform of $x(t)$.
$\overset{e}{\propto}$	The expression $f(x) \overset{e}{\propto} g(x)$ denotes that $e^{f(x)} \propto e^{g(x)}$, i.e. $f(x) = g(x) + \text{const.}$
$\delta(\cdot)$	Dirac-delta function.
$p(x), p(x y)$	For (conditional) probability density functions (PDFs), $p(x)$ and $p(x y)$ are used as shorthand notation for $p_x(x)$ and $p_{x y}(x y)$ as long as it is clear from the arguments and/or context which probability density function (PDF) the function $p(\cdot)$ refers to.
$\langle f(x) \rangle_{q(x)}$	The expression $\langle f(x) \rangle_{q(x)}$ denotes the expectation of the function $f(x)$ with respect to the random variable x distributed with PDF $q(x)$, i.e. $\langle f(x) \rangle_{q(x)} = \int f(x)q(x) \text{d}x$.
$\mathcal{D}_{\text{KL}}(q\ p)$	Kullback–Leibler divergence of q from p .

$N(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$	PDF of a real-valued normal distribution with mean $\boldsymbol{\mu} \in \mathbb{R}^N$ and covariance $\boldsymbol{\Sigma} \in \mathbb{R}^{N \times N}$, i.e. $N(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \left(\frac{1}{2\pi}\right)^{\frac{N}{2}} \boldsymbol{\Sigma} ^{-\frac{N}{2}} \times e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})}$.
$CN(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$	PDF of a complex-valued normal distribution with mean value $\boldsymbol{\mu} \in \mathbb{C}^N$, covariance $\boldsymbol{\Sigma} \in \mathbb{C}^{N \times N}$ and pseudo-covariance $\langle (\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T \rangle_{CN(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})} = \mathbf{0}$, i.e. $CN(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \left(\frac{1}{\pi}\right)^N \boldsymbol{\Sigma} ^{-N} \times e^{-(\mathbf{x}-\boldsymbol{\mu})^H \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})}$.
$Ga(x; \alpha, \beta)$	PDF of a Gamma-distributed random variable with shape α and rate β , i.e. $Ga(x; \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x}$, where $\Gamma(\cdot)$ denotes the Gamma-function.

Part I

**Introduction and Main
Results**

1

Introduction

Ultra-wideband (UWB) radios are getting more and more ubiquitous due to their increasing integration into new smartphones and cars [2]. A transmission system is considered a UWB system if it has a bandwidth of at least 500 MHz or, a relative bandwidth of one fifth of the carrier frequency [3, 4]. Due to the large bandwidth, UWB systems possess a high temporal resolution, enabling the estimation of the arrival-time of signals with nanosecond accuracy. This allows to distinguish different propagation paths in indoor environments based on the propagation delay, e.g. the separation of the direct component and reflections from the floor, walls and other objects present in the environment [3]. Hence, UWB radios can be utilized effectively to detect and estimate the position of objects and/or people in indoor environments. Due to the limited power allowed by regulators, the range of UWB radios is typically restricted to a few meters or tens of meters [2], inhibiting the use of UWB radios in long-range outdoor scenarios.

One use case for UWB radio detection and ranging (i.e. radar) systems is detecting the presence (and possibly activity) of people in indoor environments, e.g. in smart buildings, healthcare applications and car occupancy detection [5–12]. In particular, interest in car occupancy detection surged in recent years due to its inclusion in the development roadmap of the European car safety association EuroNCAP [13]. Several sensors, such as cameras, weight sensors, infrared, frequency-modulated continuous-wave (FMCW) radars or UWB radars, can be used to perform this task [11, 12]. Since UWB radios are placed in cars for keyless access systems anyways, they can provide this sensing modality in an opportunistic fashion without the additional cost and complexity of a dedicated sensor. However, multipath propagation inside the passenger cabin of the car causes time-dispersion of the signal reflected off the person. This dispersion must be accounted for in the detection algorithm.

In conventional radar systems, a widely used model is to represent each target by a single scattering point reflecting the signal [14, 15]. In such a model, the received signal consists of multiple delayed and attenuated copies of the transmitted pulse, one for each target. Transformed into the frequency domain, estimating the

propagation delays is an instance of line spectral estimation (LSE), where each target corresponds to a spectral line with (pseudo) frequency related to the targets delay. If the size of a target is larger than the range resolution determined by the temporal resolution of the system, we can observe multiple reflected copies of the transmit pulse corresponding to different parts of a single target, i.e. we observe time-dispersion of the reflected signal. In this case, the target is called an extended object [16]. The line spectrum corresponding to such an extended target consists of a cluster of closely-spaced and correlated lines. Due to this clustered structure, we refer to such a spectrum as a “structured” line spectrum. For UWB radios, a human body and other objects with similar size, can be considered as extended objects, due to the large temporal resolution of UWB systems. However, most approaches for LSE do not consider any correlation between the spectral lines.

1.1 Research Hypotheses

1.1.1 Car Occupancy Detection Using Ultra-Wideband Radar

The small space enclosed by the passenger cabin of the car results in many densely-spaced multipath components (MPCs) [17–20]. Many of those MPCs will arrive with a similar delay than the reflection from the actual target. Therefore, detecting the target and distinguishing it from a static object, e.g. a suitcase or a bag of groceries, is virtually impossible from a single measurement. However, the passenger cabin of the car can be assumed stationary, while there will always be some movements in case there is a person in the car, e.g. due to respiration. Therefore, the target can be detected based on the variation of the received signal over several snapshots [6, 11].

Furthermore, from all the MPCs inside the car, some MPCs corresponding to higher-order reflections will also interact with the target. These MPCs change over time due to the target movement. Current approaches either do not consider the MPCs corresponding to higher-order reflections interacting with the target in the detection, or add them incoherently [6, 11]. It is expected that the detection performance can be improved if these components can be added in a coherent fashion. To do so, an estimation of the dispersive multipath channel from the transmitter to the target and back to the receiver is required. Therefore, the first hypothesis investigated in this thesis is:

Hypothesis 1

The performance of car occupancy detection based on UWB radar can be improved by modeling and estimating the dispersive multipath channel between the transmitter, the target, and the receiver.

1.1.2 Estimation of Structured Line Spectra

Many approaches for LSE are based on sparsity, i.e. the problem of estimating the number of lines in the spectrum is transformed to a sparse recovery problem where the signal is represented as the product of an over-complete dictionary matrix with a sparse amplitude vector [21–26]. Most LSE algorithms assume the lines in the spectrum to be uncorrelated. Therefore, they fail account for the

correlation between the spectral lines in a structured spectrum, e.g. the spectrum of the dispersed backscatter signal from an extended object. Therefore, the second hypothesis investigated in this thesis is:

Hypothesis 2

Group sparsity can be used to generalize sparsity-based LSE methods to the case of structured line spectra. Thereby, the performance of LSE algorithms can be improved for applications where such structured line spectra arise naturally.

1.2 Contribution and Outlook

This thesis is divided into two parts. Part I introduces the research topics and summarizes the main results of the publications [A-E]. Part II is composed of the publications [A-E], which are listed below. To clarify the bibliography, the publications of Part II are referred to as [A-E] instead of the ongoing numbers used otherwise.

The remaining sections of Part I are structured as follows. Chapter 2 presents the background on which the thesis is built, introducing a model of a radio system, LSE, and variational inference (VI). Car occupancy detection using UWB radar is considered in Chapter 3 and publications [A, B, C]. Specifically, [A, B] model the radio channel from the transmitting antenna to the target and back to the receiving antenna as a backscatter channel and develop estimators for this model. In [C], we propose a data-driven occupancy detector based on the popular “ResNet” architecture. To improve the performance of detection and estimation of dispersed target signals, Chapter 4 and publications [D, E] present an extension of sparsity-based LSE approaches to structured line spectra. Specifically, in [D] we derive a fast update rule for variational block sparse Bayesian learning (SBL) that improves the algorithm runtime by two orders of magnitude compared to the block SBL (BSBL) algorithm [27, 28]. In [E] we present a generalization of super-resolution SBL to structured line spectra based on the variational expectation-maximization (EM) algorithm. The algorithm presented in [E] outperforms state-of-the-art super-resolution SBL algorithms for canonical (i.e., unstructured) line spectra in the detection of extended objects. Multi-pitch estimation from audio data is another application where structured line spectra arise naturally. The algorithm presented in [E] is shown to outperform state-of-the-art methods specifically tailored to multi-pitch estimation, demonstrating the broad applicability of the presented algorithm. Chapter 5 summarizes the main findings of this thesis and concludes Part I.

List of Included Publications

- [A] J. Möderl, F. Pernkopf, and K. Witrisal, “Car occupancy detection using UWB radar,” in *2021 18th European Radar Conference (EuRAD)*, London, U.K., Apr. 5–7, 2022, doi: 10.23919/EuRAD50154.2022.9784505
- [B] J. Möderl, E. Leitinger, F. Pernkopf, and K. Witrisal, “Variational message passing-based respiratory motion estimation and detection using radar signals,” in *2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Rhodes, Greece, Jun. 4–10, 2023, doi: ICASSP49357.2023.10095725
- [C] J. Möderl, S. Posch, F. Pernkopf, and K. Witrisal, “UWBCarGraz dataset for car occupancy detection using ultra-wideband radar,” in *2024 IEEE Radar Conference*, Denver, CO, USA, May 04–10, 2024, doi: 10.1109/Radar-Conf2458775.2024.10549280
- [D] J. Möderl, F. Pernkopf, K. Witrisal, and E. Leitinger, “Fast variational block-sparse Bayesian learning,” submitted to *IEEE Transactions on Signal Processing*, Feb. 2024, doi: 10.48550/arXiv.2306.00442
- [E] J. Möderl, F. Pernkopf, K. Witrisal, and E. Leitinger, “Variational inference of structured line spectra exploiting group-sparsity,” submitted to *IEEE Transactions on Signal Processing*, Feb. 2024, doi: 10.48550/arXiv.2303.03017

2

Background

This thesis is deeply rooted in (approximate) Bayesian statistics. Posterior probability density functions (PDFs) are central to the Bayesian approach to statistics. A posterior PDF

$$\begin{aligned} p(\mathbf{x}|\mathbf{y}) &= \frac{p(\mathbf{y}|\mathbf{x})p(\mathbf{x})}{p(\mathbf{y})} \\ &\propto p(\mathbf{y}|\mathbf{x})p(\mathbf{x}) \end{aligned} \tag{2.1}$$

represents the state of believe about the value \mathbf{x} of some random variable X after taking into consideration both data \mathbf{y} from a measurement and prior knowledge about \mathbf{x} [29,30]. In the Bayesian theorem (2.1), the evidence $p(\mathbf{y})$ models the total probability of observing \mathbf{y} over all realizations \mathbf{x} of X . The posterior $p(\mathbf{x}|\mathbf{y})$ obtained from (2.1) can be used either directly, e.g. as prior for a follow-up experiment, or by computing statistics from it, such as the mode, median, mean or higher-order moments. The mode of the posterior density is usually referred to as maximum a-posteriori (MAP) estimate.

$$\hat{\mathbf{x}}_{\text{MAP}} = \arg \max_{\mathbf{x}} p(\mathbf{x}|\mathbf{y}) = \arg \max_{\mathbf{x}} p(\mathbf{y}|\mathbf{x})p(\mathbf{x}). \tag{2.2}$$

In case no prior knowledge is considered $p(\mathbf{x}) = \text{const.}$, i.e. if all values of \mathbf{x} are considered equally likely before observing the data \mathbf{y} , the MAP estimate simplifies to the maximum likelihood (ML) estimate [30]

$$\hat{\mathbf{x}}_{\text{ML}} = \arg \max_{\mathbf{x}} p(\mathbf{y}|\mathbf{x}). \tag{2.3}$$

In practice, the Bayesian approach typically requires three ingredients: (i) a model from which the likelihood function $p(\mathbf{y}|\mathbf{x})$ is obtained, (ii) a way to encode prior knowledge, e.g. the abstract notion of sparsity, into a prior PDF $p(\mathbf{x})$, and (iii) a tractable method to compute the prior $p(\mathbf{x}|\mathbf{y})$, or an approximation thereof.

This chapter introduces the reader to the background the remaining thesis is based upon. First, a simple model of a radio system is introduced along with the multipath channel model that models the propagation environment as a (possibly infinite) sum of MPCs, which will be used as our “world model” (i). A common assumption for radar system is, that there are only a few point-like scattering objects present in the environment, i.e. the channel can be represented by a few MPCs corresponding to these objects. Estimating the number of MPCs and their respective propagation delays is an instance of LSE. Therefore, Section 2.2 introduces the LSE problem and LSE algorithms, including SBL as a sparsity-based method for LSE. In SBL, sparsity is achieved by assuming a hierarchical prior model addressing (ii). In order to relax the assumption of point-like targets to larger targets, i.e. targets that generate multiple related MPCs, the papers [D, E] expand SBL-based LSE methods to the case of structured lines spectra. The papers [B], [D] and [E] are based on VI and the variational EM algorithm. Therefore, Section 2.3 addresses (iii) by introducing VI and the variational EM algorithm as methods to approximate intractable posterior PDFs by a simpler class of proxy PDFs.

2.1 Toy Model of a Radio System

2.1.1 System Model

To obtain a model of the radio system, we consider a baseband pulse $s(t - mT_{\text{st}})$ transmitted at some multiple $m \in \mathbb{Z}$ of the pulse repetition interval T_{st} and carrier frequency f_c from some transmit antenna. After propagating through the environment, also referred to as radio channel, a receiving antenna receives the signal $r_m(t)$. The channel at time mT_{st} is modeled as linear time-invariant system¹ with impulse response $h_m(t)$, which “filters” the pulse $s(t - mT_{\text{st}})$ [31, 32]. At the receiver, a noisy version of the signal is available. The noise $w(t)$ corrupting the signal is typically modeled as additive noise, resulting in the baseband representation of the received signal [32]

$$r_m(t) = \int_{-\infty}^{\infty} s(t - mT_{\text{st}} - \xi) h_m(\xi) e^{-j2\pi f_c \xi} d\xi + w(t). \quad (2.4)$$

While this general system model is true for any radio system, Box 2.1 highlights some nuanced differences between systems intended for communications and systems intended for radar.

To continuously monitor the channel and observe changes over time, a sequence of M pulses $s_{\text{seq}}(t) = \sum_{m=0}^{M-1} s(t - mT_{\text{st}})$ is transmitted. Assuming a linear channel, the resulting received signal $r_{\text{seq}}(t)$ is a superposition of the time-shifted signals corresponding to the individual pulses

$$r_{\text{seq}}(t) = \sum_{m=0}^{M-1} \int_{-\infty}^{\infty} s(t - mT_{\text{st}} - \xi) h_m(\xi) e^{-j2\pi f_c \xi} d\xi + w(t). \quad (2.5)$$

¹Here time-invariant means invariant with respect to the time it takes for the energy of the transmitted pulse to fully dissipate, i.e. invariant over timescales of nanoseconds. If the pulse was transmitted at a different time, e.g. at $(m+1)T_{\text{st}}$, the propagation conditions might have changed resulting in a different channel $h_{m+1}(t)$.

Box 2.1: Radar and Communication Systems.

In a communication system the pulse $s(t - mT_{\text{st}})$ depends on the data and, thus, is considered unknown at the receiver. The goal is to infer which pulse $s_k(t - mT_{\text{st}})$ was transmitted out of a set of K waveforms $s_k(t)$, $k \in \{1, 2, \dots, K\}$ encoding the data to be transmitted. The channel $h_m(t)$ is considered a nuisance parameter that impairs the transmission.

In a radar system on the other hand, the pulse $s(t - mT_{\text{st}})$ is known to the receiver with the goal of inferring the channel $h_m(t)$ and, thereby, gather information about the state of the environment. E.g., the position and velocity of an inbound aircraft. However, both systems typically build upon the same radio system model (2.4) and (2.11) and differ only with respect to which variables are considered (known) parameters, (unknown) variables of interest, or (unknown) nuisance parameters.

As a more intuitive representation, the received signal $r_{\text{seq}}(t)$ in (2.5) is represented as a sequence of functions

$$r^{[m]}(\tau) \triangleq r_{\text{seq}}(mT_{\text{st}} + \tau) \quad (2.6)$$

with domain $\tau \in [0, T_{\text{st}})$ referred to as fast time. The sequence is index by a discrete variable $m \in \{0, 1, \dots, M - 1\}$ referred to as slow time. Box 2.2 further elaborates on the concepts of fast time and slow time in radar systems. Assuming a pulse $s(\tau)$ with limited duration $T_{\text{p}} \ll T_{\text{st}}$, i.e. $s(\tau) \approx 0$ for $|\tau| \geq T_{\text{p}}$, and a repetition interval T_{st} chosen large enough that the responses elicited by succeeding pulses do not interfere with each other, i.e. $h_m(\tau) \approx 0$ for $\tau < 0$ and $\tau \geq T_{\text{st}}$, then the m th element of this sequence $r^{[m]}(\tau)$ depends only on the channel $h_m(\tau)$ at slow time m but not on the channels that acted on the previously transmitted pulses

$$\begin{aligned} r^{[m]}(\tau) &\approx r_m(mT_{\text{st}} + \tau) \\ &= \int_{-\infty}^{\infty} s(\tau - \xi) h_m(\xi) e^{-j2\pi f_c \xi} d\xi + w(mT_{\text{st}} + \tau). \end{aligned} \quad (2.7)$$

For each element of the sequence $r^{[m]}(\tau)$, $N \leq \frac{T_{\text{st}}}{T_{\text{ft}}}$ samples are obtained at the receiver at fast time $\tau = nT_{\text{s}}$, $n \in \{0, 1, \dots, N - 1\}$. These samples are stacked into a vector

$$\mathbf{r}_m = [r^{[m]}(0) \ r^{[m]}(T_{\text{ft}}) \ \dots \ r^{[m]}((N - 1)T_{\text{ft}})]^T. \quad (2.8)$$

After receiving the full sequence, these vectors are either concatenated into an $NM \times 1$ vector

$$\mathbf{r} = [\mathbf{r}_1^T \ \mathbf{r}_2^T \ \dots \ \mathbf{r}_{M-1}^T]^T \quad (2.9)$$


or stacked into a $N \times M$ matrix

$$\mathbf{R} = [\mathbf{r}_1 \ \mathbf{r}_2 \ \dots \ \mathbf{r}_{M-1}] \quad (2.10)$$

such that $\mathbf{r} = \text{vec}(\mathbf{R})$. Both representations are used interchangeably depending on whichever is more convenient. Unless noted otherwise, the signal received from a single antenna is considered throughout this thesis.

Box 2.2: Slow Time and Fast Time.

From (2.6) it is obvious that the slow-time index m and fast time τ both refer to the same (absolute) time $t = mT_{\text{st}} + \tau$. By slicing the absolute time into segments with length T_{st} , a two-dimensional representation is obtained. The fast time τ represents the delay (usually in nanoseconds) relative to the transmission of the most recent pulse, i.e. the dimension over which the channel $h_m(\tau)$ acts on $s(\tau)$, whereas the slow time m indexes the evolution of the channel $h_m(\tau)$ from time $t = mT_{\text{st}}$ to $t = (m+1)T_{\text{st}}$, i.e. from one transmitted pulse to the next, as illustrated by the following figure.



This two-dimensional representation has the advantage that operations that act on the pulse, i.e. filtering each transmit pulse with some filter $g(\tau)$, act on a different (pseudo) dimension of the received signal than changes in the propagation environment, i.e. changes of h_m from the transmission of one pulse to the next. However, in this representation, the slow time m is not a continuous variable, since it can only increase in discrete steps indexing the slices of $r_{\text{seq}}(t)$ in (2.5).

This idealized model of a radio system neglects many practical effects, such as non-ideal hardware, phase and frequency synchronization, I/Q imbalance, and carrier feed-through in the demodulator. Nevertheless, it provides a useful model under the assumptions that these effects can be mitigated, e.g. by calibrating the equipment adequately.

2.1.2 Channel Model

A common model for the radio channel $h_m(\tau)$ in (2.4) is the multipath propagation model [31, 32]. This model is essentially a raytracing model, which models the propagation from the transmit to the receive antenna as a sum of K propagation

paths. For each individual propagation path, a copy of the signal is received, referred to as a MPC, that is delayed with some delay $\tau_k(t) \in \mathbb{R}_{\geq 0}$ and scaled with some amplitude $\alpha_k(t) \in \mathbb{C}$, such that

$$h_m(\tau) = \sum_{k=1}^K \alpha_k(mT_{\text{st}}) \delta(\tau - \tau_k(mT_{\text{st}})). \quad (2.11)$$

The individual MPCs represent scattering from point-like objects in the environment or specular reflections from large and flat surfaces (relative to the wavelength $\frac{f_c}{c}$, where c is the propagation speed of the signal). Changes in the propagation environment are represented by changes of the parameters $\tau_k(t)$ and $\alpha_k(t)$ over time t . As the number of MPCs tends to infinity, any propagation (through a linear medium) can be represented as a multipath channel [31, 32]. Any practical system will be band limited and for each step in slow time only N samples are obtained at the receiver. For such systems, any channel can be sufficiently represented with a finite number of $K \leq N$ MPCs. As detailed in Box 2.4 in Section 2.2, determining the number of MPCs K as well as their respective propagation delays τ_k from the received signal \mathbf{r} is an instance of LSE.

2.2 Line Spectral Estimation (LSE)

LSE is a fundamental problem present in many domains of signal processing, including, but not limited to, radar and sonar systems, geophysical exploration, vibration monitoring, meteorology, astronomy, speech analysis, control systems, and many more [33]. Reference [33, Ch. 4] provides a good overview of parametric approaches to LSE. Furthermore, [21] provides an excellent summary of more recent literature such as sparsity based approaches, including the SBL-based approach detailed in Section 2.2.1.

The data-generating model underlying the LSE problem can be expressed as


$$y(t) = \sum_{k=1}^K \alpha_k \cdot e^{j2\pi f_k t} + w(t) \quad (2.12)$$

where the continuous signal $y(t) \in \mathbb{C}$ consists of K sinusoidal components, indexed by $k \in \{1, 2, \dots, K\}$, with amplitudes $\alpha_k \in \mathbb{C}$ and frequencies $f_k \in \mathbb{R}_{\geq 0}$ in additive noise $w(t) \in \mathbb{C}$ [33]. Typically only N samples or observations $\mathbf{y} = [y(t_1) \ y(t_2) \ \dots \ y(t_N)]^T$ at sampling instances $t \in \{t_1, t_2, \dots, t_N\}$ of the continuous signal $y(t)$ are available as data. The goal of LSE is to estimate the frequencies f_k given \mathbf{y} , while the remaining parameters, e.g. the amplitudes α_k or parameters relating to the distribution of the noise process $n(t)$, are either considered nuisance parameters or known.² Box 2.3 illustrates an example of LSE while Box 2.4 details how estimating the number of scatter points in an environment and the propagation delay of the respective MPCs in a radar system can be formulated as an instance of LSE.

²In this thesis all problems of the mathematical form (2.12) are considered as LSE problems, regardless whether \mathbf{y} contains samples of a time-domain signal or from any other domain. In the same manner, this thesis refers to the parameters f_k as frequencies regardless whether they are frequencies in time, space or any other domain.

Box 2.3: Illustration of LSE.

Consider the signal $x(t) = \alpha_1 e^{j2\pi f_1 t} + \alpha_2 e^{j2\pi f_2 t}$ consisting of two sinusoidal components with frequencies $f_1 = 0.135$ Hz and $f_2 = -0.3831$ Hz and amplitudes $\alpha_1 = 0.8e^{j0.382}$ and $\alpha_2 = 0.5e^{j1.92}$. The real part of $x(t)$ is illustrated in the following plot.



The Fourier transform of $x(t)$ is


$$X(f) = \mathcal{F}\{x(t)\} = \alpha_1 \delta(f - f_1) + \alpha_2 \delta(f - f_2) \quad (2.13)$$

where $\delta(\cdot)$ denotes the Dirac-delta. Since $X(f)$ is zero everywhere except at $f \in \{f_1, f_2\}$ where it diverges, the spectrum is said to contain two “lines” at $f_1 = 0.135$ Hz and $f_2 = -0.3831$ Hz.

Sampling of $x(t)$ can be represented by multiplying the signal $x(t)$ with a finite Dirac-train $\sum_{n=0}^{N-1} \delta(t - nT_s)$ consisting of N pulses spaced with T_s . In the frequency domain, the sampled signal $X_N(f) = \mathcal{F}\{x(t) \sum_{n=0}^{N-1} \delta(t - nT_s)\} = X(f) * D_N(f)$ can be expressed as a convolution of $X(f)$ with the Dirichlet kernel

$$D_N(f) \triangleq \mathcal{F}\left\{\sum_{n=0}^{N-1} \delta(t - nT_s)\right\} = \frac{\sin(N\pi f T_s)}{\sin(\pi f T_s)} \cdot e^{-j(N-1)\pi f T_s} \quad (2.14)$$


resulting in $X_N(f) = \alpha_1 D_N(f - f_1) + \alpha_2 D_N(f - f_2)$. The plot below illustrates $|X(f)|$ and $|X_N(f)|$ obtained from sampling $x(t)$ with $N = 50$ samples at sampling interval $T_s = 1$ s.



Thus, the parameters of any signal that is either a sum of sinusoidal components or, equivalently, a sum of scaled and shifted Dirichlet kernels, can be estimated using LSE. Note, that in most cases only a noisy version of the signal $y(t) = x(t) + w(t)$ is available, where $w(t)$ is some additive noise process.

Box 2.4: Delay Estimation as LSE.

Consider the following setup of an monostatic radar system with two targets illustrated in the figure below.



By inserting the multipath channel model (2.11) with $K = 2$ paths into (2.4), the received signal $r(\tau)$ can be modeled as


$$r(\tau) = \sum_{k=1}^K \alpha_k s(\tau - \tau_k) + w(\tau) \quad (2.15)$$

where the slow time m was omitted for brevity. I.e., $r(\tau)$ consists of K copies of the transmit pulse $s(\tau)$ in noise. Each copy relates to a target modeled by a single scatter point reflecting an attenuated copy of the signal back to the receiver. The reflected copies arrive with some propagation delay $\tau_k = \frac{2d_k}{c}$ relating to the distance d_k between the radar system and the k -th target, and the signal propagation speed c . The amplitudes α_k relate to the radar cross section of each target. Transformed to the frequency domain, the received signal is expressed as

$$R(f) = \mathcal{F}\{r(\tau)\} = \sum_{k=1}^K \alpha_k S(f) e^{-j2\pi f \tau_k} + W(f) \quad (2.16)$$

where $S(f) = \mathcal{F}\{s(\tau)\}$ is the Fourier transform of the transmit pulse and $W(f) = \mathcal{F}\{w(\tau)\}$ the Fourier transform of the noise process.

A noiseless ($w(\tau) = W(f) = 0$) example of $r(\tau)$ and $R(f)$ obtained from $K = 2$ targets with delays $\tau_1 = 33$ ns, and $\tau_2 = 40$ ns corresponding to target distances $d_1 = 5$ m and $d_2 = 6$ m, respectively, is illustrated in the plot below.



Equation (2.16) is similar to (2.12). The two differences are (i) the time and frequency domains are switched and (ii) the sinusoidal bases $e^{-j2\pi f \tau_k}$ are multiplied with the spectral envelope $S(f)$ of the transmit pulse, which is equivalent to applying a time-domain windowing function in (2.12). Both of these differences can be easily accounted for in most LSE algorithms. Thus, estimating the number of targets K and their delays τ_k from the received signal $r(\tau)$ is, essentially, an LSE problem.

Existing methods for LSE can be categorized coarsely into two categories. Methods of Category (i) assume the model order K known, whereas methods of Category (ii) estimate the model order K jointly with the frequencies f_k from the data. Notable methods of Category (i) include nonlinear least-squares estimators [33], MUSIC [34], ESPRIT [35] as well as the ML method [36, 37]. However, in most practical examples the model order is not known beforehand and, therefore, has to be estimated from the data as well. A way to extend methods of Category (i) to estimate the model order as well is to evaluate the method multiple times for a range of postulated model orders, e.g. $\hat{K} \in \{0, 1, \dots, K_{\max}\}$, where K_{\max} is the largest possible model order. From among these models, the one is chosen which achieves the best tradeoff between model complexity and data fit according to an information

criteria, such as the Akaike information criterion, the Bayesian information criterion, or minimum description length [33, 38–40]. Nevertheless, this approach to estimate the model order is computationally costly, since we have to obtain estimates of the frequencies f_k for each postulated model order $\hat{K} \in \{0, 1, \dots, K_{\max}\}$.

Many methods of Category (ii), which jointly estimate the model order along with the frequencies of the spectral lines, are based on sparsity. Instances thereof include basis pursuit denoising [23] (also known as the least absolute shrinkage and selection operator (LASSO) [41]), matching pursuit [22, 42], SBL [43–47], variational line spectral estimation (VALSE) [48], semiparametric/sparse iterative covariance-based estimation (SPICE) [49, 50], and atomic noise minimization (ATM) [51, 52]. To introduce these methods, let K_{\max} be a design parameter large enough such that $K_{\max} \geq K$. The model order in (2.12) can be artificially increased from K to K_{\max} by introducing $K_{\max} - K$ additional spectral lines with arbitrary frequency and an amplitude $\alpha_k = 0 \forall k > K$. Specifically, let $\boldsymbol{\theta} = [\theta_1 \ \theta_2 \ \dots \ \theta_{K_{\max}}]^T$ be a vector of the parameters of a parameterized dictionary matrix $\boldsymbol{\Psi}(\boldsymbol{\theta}) = [\boldsymbol{\psi}(\theta_1) \ \boldsymbol{\psi}(\theta_2) \ \dots \ \boldsymbol{\psi}(\theta_{K_{\max}})] \in \mathbb{C}^{N \times K_{\max}}$ with columns $\boldsymbol{\psi}(\theta_k) = [e^{j2\pi\theta_k t_1} \ e^{j2\pi\theta_k t_2} \ \dots \ e^{j2\pi\theta_k t_N}]^T$ to express the vector \mathbf{y} obtained from sampling the generative model (2.12) at sampling times $\mathbf{t} = [t_1 \ t_2 \ \dots \ t_N]^T$ as

$$\mathbf{y} = \boldsymbol{\Psi}(\boldsymbol{\theta})\boldsymbol{\alpha} + \mathbf{w} \quad (2.17)$$

where $\boldsymbol{\alpha} = [\alpha_1 \ \alpha_2 \ \dots \ \alpha_{K_{\max}}]^T$, and $\mathbf{w} = [w(t_1) \ w(t_2) \ \dots \ w(t_N)]^T$ contains the samples of the noise process $w(t)$. Using sparse signal reconstruction methods, a sparse estimate $\hat{\boldsymbol{\alpha}}$ of $\boldsymbol{\alpha}$ is obtained together with an estimate $\hat{\boldsymbol{\theta}}$ of $\boldsymbol{\theta}$ from the inference model (2.17). Thereby, an estimates of the model order \hat{K} is indirectly obtained as the number of nonzero elements in the estimated amplitude vector $\hat{\boldsymbol{\alpha}}$, i.e. $\hat{K} = \|\hat{\boldsymbol{\alpha}}\|_0$, where $\|\cdot\|_0$ denotes the number of non-zero elements, which is sometimes also referred to as ℓ^0 -norm. Furthermore, the set $\{f_k : k \in \{1, 2, \dots, K\}\}$ of the frequencies of the generative model is estimated as the parameters θ_k of the columns of the dictionary matrix corresponding to the nonzero elements of the amplitude vector, i.e. as the set $\{\theta_k : k \in \{1, 2, \dots, K_{\max}\}, \alpha_k \neq 0\}$. Many algorithms further simplify the problem by restricting $\boldsymbol{\theta}$ to a grid as discussed in Box 2.5.

One approach to infer $\boldsymbol{\alpha}$ and $\boldsymbol{\theta}$ from (2.17) is to solve the canonical regularized regression problem

$$(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\alpha}}) = \arg \min_{\boldsymbol{\theta}, \boldsymbol{\alpha}} \|\mathbf{y} - \boldsymbol{\Psi}(\boldsymbol{\theta})\boldsymbol{\alpha}\| + g(\boldsymbol{\alpha}) \quad (2.18)$$

where the first term $\|\mathbf{y} - \boldsymbol{\Psi}(\boldsymbol{\theta})\boldsymbol{\alpha}\|$ describes the discrepancy between the data and the model, and $g(\boldsymbol{\alpha})$ is a penalty term that drives many α_k towards zero in order to obtain a sparse result. The optimization problem of (2.18) can be viewed as Bayesian MAP estimation of $\boldsymbol{\alpha}$ and ML estimation of $\boldsymbol{\theta}$ where $\|\mathbf{y} - \boldsymbol{\Psi}(\boldsymbol{\theta})\boldsymbol{\alpha}\|$ corresponds to a (log) Gaussian likelihood and $g(\boldsymbol{\alpha})$ corresponds to a (log) prior on $\boldsymbol{\alpha}$. In case $g(\boldsymbol{\alpha}) = \sum_{k=1}^{K_{\max}} \tilde{g}(\alpha_k)$ factorizes over the amplitudes α_k , this approach is an instance of Type-I estimation [57]. Box 2.6 elaborates the difference between Type-I and Type-II Bayesian estimators. Different estimators are obtained by different priors $\tilde{g}(\alpha_k)$. Common choices include $\tilde{g}(\alpha_k) \propto |\alpha_k|^p$, $p \in (0, 2]$ [23, 24, 58] and $\tilde{g}(\alpha_k) \propto \ln(|\alpha_k| + \epsilon)$, $\epsilon \geq 0$ [25, 26].

Box 2.5: Super-resolution and Grid-based Approaches.

The joint inference of $\boldsymbol{\theta}$ and $\boldsymbol{\alpha}$ in (2.17) can be quite involved due to the nonlinear dependence on $\boldsymbol{\theta}$. To simplify the problem, many algorithms assume that the true frequencies lie on a fixed grid $\bar{\boldsymbol{\theta}} = [\bar{f}_1 \ \bar{f}_2 \ \cdots \ \bar{f}_{K_{\max}}]^T$ with K_{\max} elements. Thus, only the amplitudes $\boldsymbol{\alpha}$ have to be estimated from the data [23, 42, 49, 53, 54]. The frequencies of the generative model, i.e. f_k in (2.12), are estimated as the elements of the grid $\bar{\boldsymbol{\theta}}$ that correspond to the nonzero amplitudes. The true frequencies f_k of the generative model will not lie exactly on the grid in most practical applications. Thus, this assumption introduces modeling errors which lead to a decrease in the estimation performance [55, 56].

To differentiate between approaches that assume a fixed grid for the frequencies f_k and approaches that estimate the frequencies on the continuum, we refer to approaches of the later type as super-resolution LSE. This is a more restrictive definition of super-resolution compared to the one used in some literature, e.g. [33]. Reference [33] uses super-resolution to refer to all approaches which can resolve spectral lines with a spacing smaller than a Nyquist sample, i.e. frequencies closer than $\Delta_f = \frac{1}{NT_s}$, if \mathbf{y} consists of N samples with regular sampling interval T_s . This more lenient definition of super-resolution can be achieved by grid-based approaches as well, e.g. by selecting a grid with a spacing smaller than Δ_f .

Box 2.6: Type-I and Type-II Bayesian Estimators.

Consider a hierarchical Bayesian model of data \mathbf{y} which depends on parameters \mathbf{x} via some likelihood $p(\mathbf{y}|\mathbf{x})$. The parameters \mathbf{x} depend in turn on some hyperparameters $\boldsymbol{\gamma}$, distributed as $p(\boldsymbol{\gamma})$ via a hierarchical prior $p(\mathbf{x}|\boldsymbol{\gamma})$, resulting in the joint likelihood.

$$p(\mathbf{x}, \mathbf{y}, \boldsymbol{\gamma}) = p(\mathbf{y}|\mathbf{x})p(\mathbf{x}|\boldsymbol{\gamma})p(\boldsymbol{\gamma}) \quad (2.19)$$

Assuming the primary interest lies in estimating \mathbf{x} and $\boldsymbol{\gamma}$ is considered a nuisance parameter, the Bayesian philosophy tells us to marginalize over $\boldsymbol{\gamma}$ to obtain the (Type-I) posterior PDF

$$\begin{aligned} p_I(\mathbf{x}|\mathbf{y}) &\propto \int p(\mathbf{y}|\mathbf{x})p(\mathbf{x}|\boldsymbol{\gamma})p(\boldsymbol{\gamma}) \, d\boldsymbol{\gamma} \\ &= p(\mathbf{y}|\mathbf{x}) \int p(\mathbf{x}|\boldsymbol{\gamma})p(\boldsymbol{\gamma}) \, d\boldsymbol{\gamma}. \end{aligned} \quad (2.20)$$

This direct approach to solving hierarchical models is referred to as Type-I estimation [59]. The result of the marginalization can be viewed as non-hierarchical prior $p(\mathbf{x}) = \int p(\mathbf{x}|\boldsymbol{\gamma})p(\boldsymbol{\gamma}) \, d\boldsymbol{\gamma}$. On the other hand, Type-II estimation is a two-step approach which first obtains point estimates $\hat{\boldsymbol{\gamma}}$ of the hyperparameters $\boldsymbol{\gamma}$, e.g. as MAP estimates, by marginalizing over the amplitudes

$$\begin{aligned} \hat{\boldsymbol{\gamma}} &= \arg \max_{\boldsymbol{\gamma} \succ \mathbf{0}} p(\boldsymbol{\gamma}|\mathbf{y}) \\ &= \arg \max_{\boldsymbol{\gamma} \succ \mathbf{0}} p(\boldsymbol{\gamma}) \int p(\mathbf{y}|\mathbf{x})p(\mathbf{x}|\boldsymbol{\gamma}) \, d\mathbf{x}. \end{aligned} \quad (2.21)$$

Second, the (Type-II) posterior of amplitudes is obtained as the posterior distribution conditioned on the point estimate $\hat{\boldsymbol{\gamma}}$, i.e.

$$p_{II}(\mathbf{x}|\mathbf{y}) \propto p(\mathbf{y}|\mathbf{x})p(\mathbf{x}|\hat{\boldsymbol{\gamma}}). \quad (2.22)$$

2.2.1 Sparse Bayesian Learning (SBL)

SBL [44–46, 53, 60] is a sparse signal reconstruction method that models the amplitudes using a hierarchical Bayesian model. SBL is closely related to the evidence maximization framework [61], and automatic relevance detection [62]. SBL falls into the category of Type-II estimation [59]. Note, that [57] shows Type-II estimation to include a more general class of priors $g(\boldsymbol{\alpha})$ than Type-I estimation. In SBL, the prior distribution of the amplitudes α_k are modeled to depend on hyperparameters $\boldsymbol{\gamma} = [\gamma_1 \ \gamma_2 \ \cdots \ \gamma_{K_{\max}}]^T \in \mathbb{R}_{>0}^{K_{\max}}$ as

$$p(\boldsymbol{\alpha}|\boldsymbol{\gamma}) = \prod_{k=1}^{K_{\max}} \text{CN}(\alpha_k; 0, \gamma_k^{-1}) \quad (2.23)$$

i.e. the prior for each amplitude is modeled as an independent circular-symmetric zero-mean complex Gaussian distribution with unknown precision γ_k .

First, point estimates $\hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\gamma}} = [\hat{\gamma}_1 \ \hat{\gamma}_2 \ \cdots \ \hat{\gamma}_{K_{\max}}]^T$ of the parameters $\boldsymbol{\theta}$ and hyperparameters $\boldsymbol{\gamma}$ are obtained by maximizing the (log) marginal likelihood

$$(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\gamma}}) = \arg \max_{\boldsymbol{\theta}, \boldsymbol{\gamma} > 0} \ln \int p(\mathbf{y}|\boldsymbol{\alpha}, \boldsymbol{\theta}) p(\boldsymbol{\alpha}|\boldsymbol{\gamma}) d\boldsymbol{\alpha}. \quad (2.24)$$

Equation (2.24) results in ML estimates $\hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\gamma}}$, whereas MAP estimates can be obtained by introducing prior PDFs $p(\boldsymbol{\gamma})$ and $p(\boldsymbol{\theta})$ in (2.24). In case the likelihood is Gaussian $p(\mathbf{y}|\boldsymbol{\alpha}, \boldsymbol{\theta}) = \text{CN}(\mathbf{y}; \boldsymbol{\Psi}(\boldsymbol{\theta})\boldsymbol{\alpha}, \lambda^{-1}\mathbf{I})$ with noise precision λ , the integral in (2.24) can be solved analytically leading to the SBL cost function

$$(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\gamma}}) = \arg \max_{\boldsymbol{\theta}, \boldsymbol{\gamma} > 0} \lambda^2 \mathbf{y}^H \boldsymbol{\Psi}(\boldsymbol{\theta}) \boldsymbol{\Sigma} \boldsymbol{\Psi}(\boldsymbol{\theta})^H \mathbf{y} + \ln |\boldsymbol{\Sigma}| + \sum_{k=1}^{K_{\max}} \ln \gamma_k \quad (2.25)$$

where $\boldsymbol{\Sigma} = (\lambda \boldsymbol{\Psi}(\boldsymbol{\theta})^H \boldsymbol{\Psi}(\boldsymbol{\theta}) + \boldsymbol{\Gamma})^{-1}$ and $\boldsymbol{\Gamma} = \text{diag}(\boldsymbol{\gamma})$ is a diagonal matrix with the elements of $\boldsymbol{\gamma}$ along its main diagonal. Second, the posterior distribution

$$p(\boldsymbol{\alpha}|\mathbf{y}, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\gamma}}) \propto p(\mathbf{y}|\boldsymbol{\alpha}, \hat{\boldsymbol{\theta}}) p(\boldsymbol{\alpha}|\hat{\boldsymbol{\gamma}}) \quad (2.26)$$

of the amplitudes conditional on the point estimates $\hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\gamma}}$ is evaluated to obtain an estimate $\hat{\boldsymbol{\alpha}}$ of $\boldsymbol{\alpha}$, e.g. as the expectation or mode of $p(\boldsymbol{\alpha}|\mathbf{y}, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\gamma}})$. In the case of a Gaussian likelihood, the posterior distribution of the amplitudes can be obtained analytical as

$$p(\boldsymbol{\alpha}|\mathbf{y}, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\gamma}}) = \text{CN}(\boldsymbol{\alpha}; \hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\Sigma}}) \quad (2.27)$$

with covariance $\hat{\boldsymbol{\Sigma}} = (\hat{\boldsymbol{\Gamma}} + \lambda \hat{\boldsymbol{\Psi}}^H \hat{\boldsymbol{\Psi}})^{-1}$ and mean $\hat{\boldsymbol{\alpha}} = \lambda \hat{\boldsymbol{\Sigma}} \hat{\boldsymbol{\Psi}}^H \mathbf{y}$, where $\hat{\boldsymbol{\Gamma}} = \text{diag}(\hat{\boldsymbol{\gamma}})$ and $\hat{\boldsymbol{\Psi}} = \boldsymbol{\Psi}(\hat{\boldsymbol{\theta}})$. Note, that the SBL cost function (2.25) can also be derived without assuming a hierarchical model, as described in Box 2.7. Another noteworthy approach to SBL uses VI to estimate (approximate) posterior distributions $q(\gamma_k)$ instead of point estimates $\hat{\gamma}_k$. This variational variant of SBL is introduced briefly in Box 2.8.

Box 2.7: Alternative View of SBL.

Instead of starting from the hierarchical model, SBL can also be derived from a prior $p(\boldsymbol{\alpha}) = \sup_{\boldsymbol{\gamma} \succ 0} \prod_{k=1}^{K_{\max}} \text{CN}(\alpha_k; 0, \gamma_k^{-1}) f(\gamma_k)$, where $f(\gamma_k)$ is a non-negative function [45, 57, 63, 64]. If the maximization over the latent variables $\boldsymbol{\gamma}$ is dropped, a lower bound on this prior is obtained

$$p(\boldsymbol{\alpha}) \geq \tilde{p}(\boldsymbol{\alpha}; \boldsymbol{\gamma}) = \prod_{k=1}^{K_{\max}} \text{CN}(\alpha_k; 0, \gamma_k^{-1}) f(\gamma_k). \quad (2.28)$$

Following the Type-II approach, the values of the latent variables $\boldsymbol{\gamma}$ are estimated as the values which best approximate the true prior $p(\boldsymbol{\alpha})$ in the region where the likelihood $p(\mathbf{y}|\boldsymbol{\alpha})$ is large

$$\begin{aligned} (\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\gamma}}) &= \arg \max_{\boldsymbol{\theta}, \boldsymbol{\gamma} \succ 0} \int p(\mathbf{y}|\boldsymbol{\alpha}, \boldsymbol{\theta}) |p(\boldsymbol{\alpha}) - \tilde{p}(\boldsymbol{\alpha}; \boldsymbol{\gamma})| d\boldsymbol{\alpha} \\ &= \arg \max_{\boldsymbol{\theta}, \boldsymbol{\gamma} \succ 0} \int p(\mathbf{y}|\boldsymbol{\alpha}, \boldsymbol{\theta}) \tilde{p}(\boldsymbol{\alpha}; \boldsymbol{\gamma}) d\boldsymbol{\alpha}. \end{aligned} \quad (2.29)$$

The arg max is invariant under strictly increasing transformations such as applying the logarithm. Thus, for $f(\gamma_k) = \text{const.}$, the convex approach (2.29) results in the same algorithm as maximizing the marginal likelihood (2.24) in case SBL is derived via the hierarchical model of (2.23).

The convex approach to SBL has advantages in certain cases, since it can represent a more general class of priors than the ones obtained by the hierarchical model [64]. Nevertheless, if a corresponding hierarchical model exists, it results in the same algorithm [64]. In this thesis, the differences between the two approaches are considered mostly of semantic nature, whether the variables $\boldsymbol{\gamma}$ are considered random variables of an underlying hierarchical model or latent variables of a convex dual-form representation of the prior $p(\boldsymbol{\alpha})$.

Box 2.8: Brief Introduction to Variational SBL.

Instead of directly maximizing the marginal integral (2.24), the maximization is sometimes performed using the EM algorithm, e.g. in the original work by Tipping [62]. A natural extension thereof is to estimate a posterior distribution also for the hyperparameters using VI. Considering a fixed dictionary $\boldsymbol{\Psi}(\bar{\boldsymbol{\theta}})$ (see Box 2.5), [53] introduces a prior distribution $p(\boldsymbol{\gamma})$ on the hyperparameters and applies VI to approximate the posterior distribution

$$p(\boldsymbol{\alpha}, \boldsymbol{\gamma}|\mathbf{y}) \propto p(\mathbf{y}|\boldsymbol{\alpha}) p(\boldsymbol{\alpha}|\boldsymbol{\gamma}) p(\boldsymbol{\gamma}) \quad (2.30)$$

with a factorized distribution

$$q(\boldsymbol{\alpha}, \boldsymbol{\gamma}) = q_{\boldsymbol{\alpha}}(\boldsymbol{\alpha}) \prod_{k=1}^{K_{\max}} q_{\boldsymbol{\gamma}, k}(\gamma_k) \quad (2.31)$$

using the variational mean-field approximation detailed in Section 2.3.2. Due to the interdependence between the weights $\boldsymbol{\alpha}$ and the hyperparameters $\boldsymbol{\gamma}$, the VI algorithm exhibits slow convergence. However, a fast update accelerating convergence was presented in [53]. Furthermore, this VI approach to SBL was extended

to include the estimation of dictionary parameters based on the variational-EM algorithm in [47]. The VALSE algorithm [48] continues this line of research and estimates approximate posterior densities also for the dictionary parameters θ . However, it is based on a slightly different hierarchical model that uses binary existence variables as hyperparameters instead of the precision hyperparameters used in SBL. The interested reader is referred to [57, 64] for a more in-depth discussion on the similarities and differences between the marginalized likelihood, convex, and VI approach to SBL.

2.3 Variational Inference (VI)

This introduction of VI is based on [65, 66] as well as on the excellent textbook by C. M. Bishop [29]. For further details, the reader is referred to these references.

2.3.1 The Variational Bayesian Approach

Computing and evaluating the posterior PDF (2.1) is intractable in many practical problems [29, Ch. 10]. Thus, many methods have been developed to approximate intractable posterior PDF. One approach is to approximate the posterior PDF with a simpler distribution. Consider a model with observed data \mathbf{y} and parameters \mathbf{x} , also referred to as hidden or latent variables in some communities. Given some approximating PDF $q(\mathbf{x}) \in \mathcal{Q}$ that is a member of some class of distributions \mathcal{Q} , the log evidence $\ln p(\mathbf{y})$ can be decomposed into

$$\ln p(\mathbf{y}) = \mathcal{L}(q) + \mathcal{D}_{\text{KL}}(q \| p_{\mathbf{x}|\mathbf{y}}) \quad (2.32)$$

where

$$\mathcal{D}_{\text{KL}}(q \| p_{\mathbf{x}|\mathbf{y}}) = \left\langle \ln \frac{q(\mathbf{x})}{p(\mathbf{x}|\mathbf{y})} \right\rangle_{q(\mathbf{x})} = - \left\langle \ln \frac{p(\mathbf{x}|\mathbf{y})}{q(\mathbf{x})} \right\rangle_{q(\mathbf{x})} \quad (2.33)$$

is the Kullback-Leibler (KL) divergence from the approximating distribution $q(\mathbf{x})$ to the posterior distribution $p(\mathbf{x}|\mathbf{y})$ and

$$\mathcal{L}(q) = \left\langle \ln p(\mathbf{x}, \mathbf{y}) - \ln q(\mathbf{x}) \right\rangle_{q(\mathbf{x})} \quad (2.34)$$

is a functional of the approximating distribution $q(\mathbf{x})$ [66], [29, Ch. 10].

The decomposition of the evidence (2.32) is proven by computing the expectation of $\ln p(\mathbf{y})$ over \mathbf{x} being distributed with PDF $q(\mathbf{x})$

$$\begin{aligned} \ln p(\mathbf{y}) &= \left\langle \ln p(\mathbf{y}) \right\rangle_{q(\mathbf{x})} \\ &= \left\langle \ln \left[p(\mathbf{y}) \frac{p(\mathbf{x}|\mathbf{y}) q(\mathbf{x})}{p(\mathbf{x}|\mathbf{y}) q(\mathbf{x})} \right] \right\rangle_{q(\mathbf{x})} \\ &= \left\langle \ln \frac{p(\mathbf{x}, \mathbf{y})}{q(\mathbf{x})} \right\rangle_{q(\mathbf{x})} + \left\langle \ln \frac{q(\mathbf{x})}{p(\mathbf{x}|\mathbf{y})} \right\rangle_{q(\mathbf{x})} \\ &= \mathcal{L}(q) + \mathcal{D}_{\text{KL}}(q \| p_{\mathbf{x}|\mathbf{y}}). \end{aligned} \quad (2.35)$$

If it exists, the KL divergence $\mathcal{D}_{\text{KL}}(q \| p_{\mathbf{x}|\mathbf{y}}) \geq 0$ is a nonnegative quantity which is zero if, and only if, $q(\mathbf{x}) = p(\mathbf{x}|\mathbf{y})$. From the decomposition of the evidence (2.32), it follows that $\ln p(\mathbf{y}) \geq \mathcal{L}(q)$, i.e. $\mathcal{L}(q)$ is a lower bound on the model evidence,

referred to as the evidence lower bound (ELBO). Because the model evidence $p(\mathbf{y})$ is a fixed quantity, we can minimize the KL divergence $\mathcal{D}_{\text{KL}}(q||p_{\mathbf{x}|\mathbf{y}})$ by maximizing the ELBO $\mathcal{L}(q)$. However, when maximizing the ELBO

$$q^* = \arg \max_{q \in \mathcal{Q}} \mathcal{L}(q) \quad (2.36)$$

without restricting the solution space \mathcal{Q} , the maximum is clearly achieved when $\mathcal{D}_{\text{KL}}(q||p_{\mathbf{x}|\mathbf{y}}) = 0$ which is the case exactly if $q(\mathbf{x}) = p(\mathbf{x}|\mathbf{y})$. I.e. finding the distribution q^* amounts to finding the true posterior $p(\mathbf{x}|\mathbf{y})$ which is intractable by the assumption motivating the VI approach in the first place. Hence, the VI approach seems to offer no advantage compared to finding and evaluating the posterior $p(\mathbf{x}|\mathbf{y})$ directly using (2.1).

Nevertheless, the decomposition of the evidence $\ln p(\mathbf{y})$ is useful since maximizing the ELBO provides a systematic way of finding an approximating PDF $q(\mathbf{x})$ that is closest to the true posterior $p(\mathbf{x}|\mathbf{y})$ in terms of the KL divergence. The optimization problem in (2.36) can be simplified by constraining the space of approximating distributions $q \in \mathcal{Q}$ to some restricted space \mathcal{Q} . If a very simple class of approximating distributions is chosen, e.g. a parameterized family of distributions $q(\mathbf{x}; \boldsymbol{\omega})$ with parameters $\boldsymbol{\omega}$, then the ELBO $\mathcal{L}(q)$ becomes a function of those parameters. In this case, the infinite-dimensional optimization of (2.36) is reduced to optimizing the ELBO over the parameter vector $\boldsymbol{\omega}$. The posteriors $p(\mathbf{x}|\mathbf{y})$ will typically not be a member of this simple class of distributions, resulting in an approximation error $\mathcal{D}_{\text{KL}}(q||p_{\mathbf{x}|\mathbf{y}}) > 0$ and, equivalently, $\mathcal{L}(q) < \ln p(\mathbf{y})$ being strictly smaller than $\ln p(\mathbf{y})$. On the other hand, if a more expressive class of distributions is chosen for \mathcal{Q} , this will typically result in a better approximation, i.e. smaller KL divergence $\mathcal{D}_{\text{KL}}(q||p_{\mathbf{x}|\mathbf{y}})$. However, a more expressive class of distributions will typically require more parameters to describe, increasing the complexity of the optimization of the ELBO in (2.36). Thus, the selection of the solution class \mathcal{Q} can be used to achieve a tradeoff between better approximation quality (in terms of the KL divergence) and the complexity of the optimization problem in (2.36).

2.3.2 Mean-Field Approximation

A common approximation to the posterior is the so called mean-field approximation [65, 66], [29, Ch. 10]. Let \mathbf{x} be partitioned into M disjoint sub vectors \mathbf{x}_i , $i \in \{1, 2, \dots, M\}$, such that $\mathbf{x} = [\mathbf{x}_1^T \mathbf{x}_2^T \dots \mathbf{x}_M^T]^T$. The mean-field approximation restricts \mathcal{Q} to be the set of all factorized distributions of the form

$$q(\mathbf{x}) = \prod_{i=1}^M q_i(\mathbf{x}_i) \quad (2.37)$$

where $q_i(\mathbf{x}_i)$ is an arbitrary PDF of \mathbf{x}_i . To maximize the ELBO, we now seek to make a free form optimization of $\mathcal{L}(q)$ with respect to all of the distributions $q_i(\mathbf{x}_i)$. Let q_i be a shorthand notation for $q_i(\mathbf{x}_i)$, the dependence of the ELBO on one

factor q_i of $q(\mathbf{x})$ can be expressed by inserting (2.37) into (2.34) as

$$\begin{aligned}
\mathcal{L}(q) &= \left\langle \ln p(\mathbf{x}, \mathbf{y}) - \sum_{i=1}^M \ln q_i \right\rangle_{\prod_{i=1}^M q_i} \\
&= \left\langle \left\langle \ln p(\mathbf{x}, \mathbf{y}) - \ln q_i - \sum_{j=1, j \neq i}^M \ln q_j \right\rangle_{\prod_{j \neq i} q_j} \right\rangle_{q_i} \\
&= \left\langle \ln \tilde{p}_i(\mathbf{x}_i, \mathbf{y}) - \ln q_i \right\rangle_{q_i} - \underbrace{\sum_{j=1, j \neq i}^M \left\langle \ln q_j \right\rangle_{q_j}}_{\text{const. w.r.t. } q_i} \\
&= -\mathcal{D}_{\text{KL}}(q_i \| \tilde{p}_i) + \text{const.}
\end{aligned} \tag{2.38}$$

where the distribution $\tilde{p}_i(\mathbf{x}_i, \mathbf{y})$ was defined via the relation

$$\ln \tilde{p}_i(\mathbf{x}_i, \mathbf{y}) = \left\langle \ln p(\mathbf{x}, \mathbf{y}) \right\rangle_{\prod_{j \neq i} q_j} - \ln Z_i \tag{2.39}$$

and $Z_i = \int \exp \left\langle \ln p(\mathbf{x}, \mathbf{y}) \right\rangle_{\prod_{j \neq i} q_j} d\mathbf{x}_i$ is a normalization constant such that \tilde{p}_i is a valid PDF that integrates to 1. From (2.38) it follows the ELBO is maximized with respect to one of the factors q_i of $q(\mathbf{x})$ by minimizing the KL divergence $\mathcal{D}_{\text{KL}}(q_i \| \tilde{p}_i)$ from q_i to \tilde{p}_i . The minimum of the KL divergence $\mathcal{D}_{\text{KL}}(q_i \| \tilde{p}_i) = 0$ is achieved if, and only if, $q_i(\mathbf{x}_i) = \tilde{p}_i(\mathbf{x}_i, \mathbf{y})$. Thus, considering the remaining factors q_j , $j \neq i$ fixed, the factor q_i that maximizes the ELBO in (2.36) is

$$q_i(\mathbf{x}_i) = \frac{1}{Z_i} \exp \left\langle \ln p(\mathbf{x}, \mathbf{y}) \right\rangle_{\prod_{j \neq i} q_j}. \tag{2.40}$$

In practice, the normalization constant Z_i is often irrelevant and obtained only when explicitly needed, resulting in the well known factor update expression

$$\ln q_i(\mathbf{x}_i) \stackrel{e}{\propto} \left\langle \ln p(\mathbf{x}, \mathbf{y}) \right\rangle_{\prod_{j \neq i} q_j} \tag{2.41}$$

where $\stackrel{e}{\propto}$ is used to denote that the left side is equal to the right side except for an additive constant, i.e. after exponentiating both sides they are proportional to each other. Note, that through the expectation in (2.40) the optimal factor q_i depends on other factors q_j , $j \neq i$, of q . Hence, the ELBO is iteratively maximized with respect to one factor after the other until convergence.

In many cases, the mean-field approximation together with conjugate priors (see Box 2.9) leads to update-equations (2.40) that result in distributions from a specific parameterized family, e.g. the family of Gaussian distributions [65]. One such example is detailed in Box 2.10. In these cases, each update amounts to updating the parameters, e.g. the mean and covariance for a Gaussian distribution, which can often be done analytically. This results in an interdependent set of update equations for the parameters of all factors q_i of $q(\mathbf{x})$. In case the joint distribution is represented as a graphical model, e.g. using factor graphs, the iterative algorithm resulting from (2.40) can be expressed as a message passing algorithm on this factor graph [67].

Box 2.9: Conjugate Priors.

A prior $p(x)$ is conjugate to some likelihood $p(y|x)$ if the posterior $p(x|y) \propto p(y|x)p(x)$ is in the same family of distributions as the prior $p(x)$. Conjugate priors are an algebraic convenience, since they allow for closed-form expressions of the posterior $p(x|y)$. For example, the conjugate priors for the (unknown) mean x of a Gaussian likelihood $p(y|x) = \mathcal{N}(y; x, \sigma^2)$ with known variance σ^2 , is a Gaussian distribution $p(x) = \mathcal{N}(x; x_0, \sigma_0^2)$. The resulting posterior $p(x|y) = \mathcal{N}(x; \hat{x}, \hat{\sigma}^2)$ is Gaussian distributed with mean \hat{x} and variance $\hat{\sigma}^2$ given by

$$\hat{x} = \hat{\sigma}^2 \left(\frac{x_0}{\sigma_0^2} + \frac{y}{\sigma^2} \right) \quad \hat{\sigma}^2 = \left(\frac{1}{\sigma_0^2} + \frac{1}{\sigma^2} \right)^{-1}. \quad (2.42)$$

Box 2.10: Variational Mean-Field for Linear Regression.

Consider the following linear model

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{w} \quad (2.43)$$

in additive white Gaussian noise (AWGN) \mathbf{w} with likelihood $p(\mathbf{y}|\mathbf{x}) = \mathcal{N}(\mathbf{y}; \mathbf{A}\mathbf{x}, \lambda\mathbf{I})$. The goal is to infer the unknown coefficients $\mathbf{x} \in \mathbb{R}^N$ given the data \mathbf{y} , the matrix \mathbf{A} and noise precision λ .

Assuming a zero-mean Gaussian prior with unit variance $p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{0}, \mathbf{I})$, the posterior follows from Bayes theorem (2.1) as

$$p(\mathbf{x}|\mathbf{y}) = \mathcal{N}(\mathbf{x}; \hat{\mathbf{x}}, \boldsymbol{\Sigma}) \quad (2.44)$$

where $\hat{\mathbf{x}} = \lambda \boldsymbol{\Sigma} \mathbf{A}^T \mathbf{y}$ is the posterior mean and $\boldsymbol{\Sigma} = (\lambda \mathbf{A}^T \mathbf{A} + \mathbf{I})^{-1}$ the posterior covariance, i.e. the prior is conjugate to the likelihood. Nevertheless, for large dimensions N of \mathbf{x} , the calculation of the $N \times N$ covariance matrix $\boldsymbol{\Sigma}$ becomes intractable. To reduce the computational complexity, the vector \mathbf{x} can be partitioned into $M = \frac{N}{K}$ partitions \mathbf{x}_i of equal size K to approximate the posterior with a factorized distribution $q(\mathbf{x}) = \prod_{i=1}^M q_i(\mathbf{x}_i)$. From (2.40) the optimal factor updates q_i are obtained as

$$\begin{aligned} \ln q_i(\mathbf{x}_i) &\stackrel{e}{\propto} \left\langle \ln p(\mathbf{y}|\mathbf{x}) + \ln p(\mathbf{x}) \right\rangle_{\prod_{j \neq i} q_j} \\ &\propto \left\langle -\frac{\lambda}{2} \left(\mathbf{y} - \mathbf{A}_i \mathbf{x}_i - \sum_{j \neq i} \mathbf{A}_j \mathbf{x}_j \right)^T \left(\mathbf{y} - \mathbf{A}_i \mathbf{x}_i - \sum_{j \neq i} \mathbf{A}_j \mathbf{x}_j \right) \right. \\ &\quad \left. - \frac{1}{2} \mathbf{x}_i^T \mathbf{x}_i \right\rangle_{\prod_{j \neq i} q_j} \\ &\stackrel{e}{\propto} -\frac{1}{2} \mathbf{x}_i^T (\lambda \mathbf{A}_i^T \mathbf{A}_i + \mathbf{I}) \mathbf{x}_i - \mathbf{x}_i^T \lambda \mathbf{A}_i^T \left(\mathbf{y} - \sum_{j \neq i} \mathbf{A}_j \langle \mathbf{x}_j \rangle_{q_j} \right) \\ &\stackrel{e}{\propto} \ln \mathcal{N}(\mathbf{x}_i; \hat{\mathbf{x}}_i, \boldsymbol{\Sigma}_i) \end{aligned} \quad (2.45)$$

where the columns of the matrix $\mathbf{A} = [\mathbf{A}_1 \mathbf{A}_2 \cdots \mathbf{A}_M]$ are partitioned in the same manner as \mathbf{x} . The resulting distributions q_i are Gaussian distributions with mean $\hat{\mathbf{x}}_i = \lambda \boldsymbol{\Sigma}_i \mathbf{A}_i^T \mathbf{y}_{\text{res}}$ and covariance $\boldsymbol{\Sigma}_i = (\lambda \mathbf{A}_i^T \mathbf{A}_i + \mathbf{I})^{-1}$, where $\mathbf{y}_{\text{res}} = \mathbf{y} - \sum_{j \neq i} \mathbf{A}_j \hat{\mathbf{x}}_j$ is the residual signal after subtracting the (mean) contribution of the remaining coefficients.

Using the variational mean-field approach, it is sufficient to invert K matrices of size $M \times M$, $M = \frac{N}{K} \ll N$, which is of much lower complexity than inverting a $N \times N$ matrix for the true posterior. In the fully factorized case, i.e. $K = N$, the computational complexity is reduced from $\mathcal{O}(N^3)$ to an iterative algorithms where the calculation of the covariance in each iteration requires only a scalar division for each element of \mathbf{x} , resulting in complexity $\mathcal{O}(N)$. This computational gain incurs the cost of not being able to consider correlations between the different partitions of the coefficients \mathbf{x} .

2.3.3 Variational Expectation-Maximization (EM) Algorithm

Although the update equation (2.40) can lead to tractable and/or analytic update expressions for the factors q_i of $q(\mathbf{x})$, this is not always the case. If no analytic solution to the update expression (2.40) can be found, a nonparametric approach can be applied, e.g. based on a particle representation. Another possibility is to refrain from estimating a distribution for those variables and instead resort to point estimates, e.g. in an approach known as the variational EM algorithm [66].

Consider a joint PDF $p(\mathbf{x}, \boldsymbol{\theta}, \mathbf{y}) = p(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta})p(\mathbf{x})p(\boldsymbol{\theta})$ comprised of a likelihood $p(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta})$ depending on two sets of variables, i.e. \mathbf{x} and $\boldsymbol{\theta}$, and (independent) priors $p(\mathbf{x})$ and $p(\boldsymbol{\theta})$. Lets assume that some nontrivial dependence of $p(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta})$ on $\boldsymbol{\theta}$ leads to an intractable posterior distribution $p(\boldsymbol{\theta}, \mathbf{x}|\mathbf{y})$. Lets further assume that even if a mean-field approach $q(\mathbf{x}, \boldsymbol{\theta}) = q_{\mathbf{x}}(\mathbf{x})q_{\boldsymbol{\theta}}(\boldsymbol{\theta})$ is applied, no tractable solution for $q_{\boldsymbol{\theta}}(\boldsymbol{\theta})$ can be found. Since estimating a posterior is intractable, a point estimate $\hat{\boldsymbol{\theta}}$ of $\boldsymbol{\theta}$ is obtained by maximizing the evidence

$$\hat{\boldsymbol{\theta}} = \arg \max_{\boldsymbol{\theta}} p(\boldsymbol{\theta}, \mathbf{y}) = \arg \max_{\boldsymbol{\theta}} \int p(\mathbf{x}, \boldsymbol{\theta}, \mathbf{y}) d\mathbf{x}. \quad (2.46)$$

Box 2.11 discusses an alternative interpretation of (2.46) as the evidence $p(\mathbf{y}; \boldsymbol{\theta})$ of \mathbf{y} parameterized by deterministic but unknown parameters $\boldsymbol{\theta}$.

In case the marginalization over the variables \mathbf{x} is intractable, an approximating distribution $q(\mathbf{x}) \in \mathcal{Q}$ can be introduced to simplify this marginalization. Replacing the evidence $p(\mathbf{y})$ with $p(\boldsymbol{\theta}, \mathbf{y})$ in (2.32), $p(\boldsymbol{\theta}, \mathbf{y})$ can be decomposed into

$$\ln p(\boldsymbol{\theta}, \mathbf{y}) = \mathcal{L}(q, \boldsymbol{\theta}) + \mathcal{D}_{\text{KL}}(q \| p_{\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}}) \quad (2.47)$$

using derivations similar to (2.35), where $\mathcal{D}_{\text{KL}}(q \| p_{\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}})$ is the KL divergence from the approximating distribution $q(\mathbf{x})$ to the posterior $p(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta})$ of \mathbf{x} given \mathbf{y} and $\boldsymbol{\theta}$. In this case, the ELBO

$$\mathcal{L}(q, \boldsymbol{\theta}) = \langle \ln p(\mathbf{x}, \boldsymbol{\theta}, \mathbf{y}) - \ln q(\mathbf{x}) \rangle_{q(\mathbf{x})} \quad (2.48)$$

is not only a functional of the distribution $q(\mathbf{x})$, but also conditioned on $\boldsymbol{\theta}$ via the dependence of $\ln p(\mathbf{x}, \boldsymbol{\theta}, \mathbf{y})$ on $\boldsymbol{\theta}$. Using (2.47), point estimates $\hat{\boldsymbol{\theta}}$ as well as the approximating distribution $q(\mathbf{x})$ are found by maximizing the ELBO $\mathcal{L}(q, \boldsymbol{\theta})$ instead of maximizing (2.46)

$$(q^*, \hat{\boldsymbol{\theta}}) = \arg \max_{q, \boldsymbol{\theta}} \mathcal{L}(q, \boldsymbol{\theta}). \quad (2.49)$$

The maximization is typically performed in an iterative manner. In the expectation or E-step, we find the distribution $q_{\hat{\boldsymbol{\theta}}^{\text{old}}}(\mathbf{x})$ that maximizes the ELBO

$$q_{\hat{\boldsymbol{\theta}}^{\text{old}}}(\mathbf{x}) = \arg \max_q \mathcal{L}(q, \hat{\boldsymbol{\theta}}^{\text{old}}) \quad (2.50)$$

Box 2.11: Are θ Random Variables or Deterministic Parameters?

Equation (2.46) can be interpreted as the evidence of a statistical model $p(\mathbf{y}|\mathbf{x};\theta)$ with random variables \mathbf{x} parameterized by “deterministic but unknown” parameters θ . In fact, this description is used by many references, e.g. [29, 66]. However, in the Bayesian perspective, there are no “deterministic unknown parameters”. All unknown variables are considered random variables and assigned an associated prior PDF, i.e. a function that characterizes our belief about the value of these variables prior to performing the experiment. This Bayesian viewpoint is adopted throughout Part I of this thesis, although not necessarily throughout the papers comprising Part II. From this viewpoint, MAP estimation of θ can be considered as finding the “best” approximation to the posterior $p(\theta|\mathbf{y})$ using a Dirac-delta distribution $\delta(\theta - \hat{\theta})$. However, “best” in this case can not be defined using the KL divergence, since the KL divergence, and by extension the ELBO, are not defined for Dirac-delta distributions. This deviation from fully Bayesian inference is sometimes necessary in order to obtain tractable estimators, analogous to using the mean-field approximation to approximate an intractable posterior with a factorized distribution.

with respect to q while considering the current estimate $\hat{\theta}^{\text{old}}$ fixed. Following a similar derivation as (2.38), this distribution is found to be

$$\ln q_{\hat{\theta}^{\text{old}}}(\mathbf{x}) \stackrel{e}{\propto} \ln p(\mathbf{x}, \hat{\theta}^{\text{old}}, \mathbf{y}). \quad (2.51)$$

In the following maximization or M-step, the ELBO $\mathcal{L}(q_{\hat{\theta}^{\text{old}}}, \theta)$ is maximized with respect to θ while considering $q = q_{\hat{\theta}^{\text{old}}}$ fixed by

$$\begin{aligned} \hat{\theta}^{\text{new}} &= \arg \max_{\theta} \mathcal{L}(q_{\hat{\theta}^{\text{old}}}, \theta) \\ &= \arg \max_{\theta} \langle \ln p(\mathbf{x}, \theta, \mathbf{y}) \rangle_{q_{\hat{\theta}^{\text{old}}}}. \end{aligned} \quad (2.52)$$

Note, that the distribution $q_{\hat{\theta}^{\text{old}}}$ over which the expectation in (2.52) is taken does not depend on θ , only on its previous estimate $\hat{\theta}^{\text{old}}$. Thus, the moments of $q_{\hat{\theta}^{\text{old}}}$ needed to solve the optimization problem can be computed beforehand and, often, they can be computed analytically.

In many cases, the variational EM algorithm is applied together with a mean-field approximation restricting $q(\mathbf{x})$ to the class of factorized distributions $q(\mathbf{x}) = \prod_{i=1}^M q_i(\mathbf{x}_i)$ based on partitioning \mathbf{x} into several disjoint subsets of variables (see Section 2.3.2). In this case, the distribution $q_{i, \hat{\theta}^{\text{old}}}$ that maximizes the ELBO with respect to q_i while considering $\hat{\theta}^{\text{old}}$ as well as the remaining distributions q_j , $j \neq i$ fixed is found by the relation

$$\ln q_{i, \hat{\theta}^{\text{old}}}(\mathbf{x}_i) \stackrel{e}{\propto} \langle \ln p(\mathbf{x}, \hat{\theta}^{\text{old}}, \mathbf{y}) \rangle_{\prod_{j \neq i} q_j}. \quad (2.53)$$

The variational-EM algorithm is illustrated on a frequency estimation example in Box 2.12.

Box 2.12: Variational EM for Frequency Estimation.

Consider estimating the frequency f of a sinusoidal signal

$$\mathbf{y} = \alpha e^{j2\pi f \mathbf{t}} + \mathbf{w}, \quad \mathbf{y} \in \mathbb{C}^N \quad (2.54)$$

with unknown amplitude $\alpha \in \mathbb{C}$ in AWGN \mathbf{w} with PDF $p(\mathbf{w}) = \text{CN}(\mathbf{w}; \mathbf{0}, \lambda^{-1} \mathbf{I})$ and unknown precision λ , where $\mathbf{t} = [t_1 \ t_2 \ \dots \ t_N]^T$ is a length- N vector of the sampling times of the signal. From the AWGN assumption, the likelihood follows as

$$p(\mathbf{y}|f, \alpha, \lambda) = \text{CN}(\mathbf{y}; \alpha e^{j2\pi f \mathbf{t}}, \lambda^{-1} \mathbf{I}). \quad (2.55)$$

For the amplitude α and noise precision λ we can define conjugate priors, such that the posteriors can be calculated analytically. However, for the frequency f no conjugate prior is known. Introducing priors $p(\alpha) = \text{CN}(\alpha; 0, 1)$, $p(\lambda) \propto \lambda^{-1}$, and $p(f)$, the joint PDF is written as

$$p(\mathbf{x}, \boldsymbol{\theta}, \mathbf{y}) = p(f, \alpha, \lambda, \mathbf{y}) = p(\mathbf{y}|f, \alpha, \lambda) p(f) p(\alpha) p(\lambda). \quad (2.56)$$

To estimate f , α and λ , a variational EM approach is applied in which $\boldsymbol{\theta} = [f]$ and $\mathbf{x} = [\alpha \ \lambda]^T$, i.e. posterior distributions are estimated for α and λ , whereas a point estimate is obtained for f . For the sake of illustration, the posterior $p(\alpha, \lambda|f, \mathbf{y})$ is approximated using a mean-field approximation with a factorized distribution $q(\alpha, \lambda) = q_\alpha(\alpha) q_\lambda(\lambda)$. The approximating distributions $q_\alpha(\alpha)$ and $q_\lambda(\lambda)$ that maximize the ELBO conditional on the current estimate \hat{f} of f are obtained by inserting (2.56) in (2.53) as

$$\begin{aligned} \ln q_\alpha(\alpha) &\stackrel{e}{\propto} -((\hat{\lambda}N + 1)|\alpha|^2 - \text{Re}\{2\hat{\lambda}\mathbf{y}^H e^{-j2\pi\hat{f}\mathbf{t}}\alpha\}) \\ &\stackrel{e}{\propto} -\sigma_\alpha^2 |\alpha - \hat{\alpha}|^2 \end{aligned} \quad (2.57)$$

$$\ln q_\lambda(\lambda) \stackrel{e}{\propto} (N - 1) \ln \lambda - \lambda (\|\mathbf{y} - e^{j2\pi\hat{f}\mathbf{t}}\hat{\alpha}\|^2 + N\sigma_\alpha^2) \quad (2.58)$$

after solving the expectation, “completing the square”, and a few algebraic manipulations. I.e., q_α is a Gaussian distribution with mean $\hat{\alpha} = \hat{\lambda}\sigma_\alpha^2 \mathbf{y}^H e^{-j2\pi\hat{f}\mathbf{t}}$ and variance $\sigma_\alpha^2 = (\hat{\lambda}N + 1)^{-1}$, and $q_\lambda \propto \lambda^{N-1} e^{-\beta\lambda}$ is a Gamma distribution with shape N and rate $\beta = \|\mathbf{y} - e^{j2\pi\hat{f}\mathbf{t}}\hat{\alpha}\|^2 + N\sigma_\alpha^2$. The frequency \hat{f} can be estimated based on the current estimates of q_α and q_λ from (2.52) as

$$\hat{f} = \arg \max_f \|\mathbf{y} - e^{j2\pi f \mathbf{t}} \hat{\alpha}\|^2 + \ln p(f). \quad (2.59)$$

The ELBO is maximized by iterating the interdependent equations (2.57)–(2.59) between the estimate \hat{f} of f and the parameters $\hat{\alpha}$, σ_α^2 and $\hat{\lambda} = \frac{N}{\beta}$ describing q_α and q_λ , respectively.

3

Car Occupancy Detection

This chapter investigates Hypothesis 1, i.e. it deals with the problem of detecting whether a car is occupied or not using signals from UWB radars placed inside the passenger compartment of the car. The detection is performed once the car is parked and the driver exits the vehicle. In order to alert the driver while he is still in the vicinity of the car, the detection process should be finished within a few seconds after the driver left the vehicle [13].

Generally, the goal is to obtain a binary decision between the two statistical models

Model \mathcal{H}_0 : the car is empty, and

Model \mathcal{H}_1 : the car is occupied

based on the received signal observing the radio channel as a function of fast time (i.e. propagation delay) at difference instance in (slow) time. The received signal is represented either as fast time/slow time matrix \mathbf{R} or stacked into a long vector $\mathbf{r} = \text{vec}(\mathbf{R})$, as described in Section 2.1.

Hypothesis 1 postulates that considering the multipath channel in the detection algorithm leads to a better detection performance. Section 3.1 derives a model for the case of small target movements and details the model-based detection algorithms published in [A,B]. However, deriving a useful model for large movements is not trivial. To detect large target movements, Section 3.2 introduces the UWBCarGraz dataset [68] and a data-driven approach for car occupancy detection published in [C], which is able to cope with both large and small target movements. Section 3.3 compares the performance of the developed detection algorithms and Section 3.4 concludes the chapter with a short discussion of the obtained results.

3.1 Model-driven Car Occupancy Detection

3.1.1 Model of the Empty Car

To develop model-driven car occupancy detection algorithms, first a model of the received signal \mathbf{r} under both statistical models \mathcal{H}_0 and \mathcal{H}_1 is needed. To obtain a likelihood for \mathbf{r} under \mathcal{H}_0 , i.e. for the empty car, we start with the general model of a radio system introduced in Section 2.1. Without loss of generality, the sum in the multipath channel model (2.11) can be split up into L MPCs that interact with the target and $K - L$ MPCs that do not interact with the target. Since the car is assumed stationary, all propagation paths inside the passenger compartment that do not interact with the target are constant as well. For the environment outside the passenger cabin of the car, this assumptions might not be true, e.g. due to people or other cars moving around. However, all propagation paths to these moving objects have to pass through the enclosing of the passenger compartment twice, resulting in a high attenuation of these propagation paths. Hence, their contribution to the channel is considered irrelevant. Assumption 1 defines this formally.

Assumption 1. *For $k > L$, i.e. for MPCs that do not interact with the target, the parameters $\alpha_k(mT_{\text{st}}) = \alpha_k$ and $\tau_k(mT_{\text{st}}) = \tau_k$, are constant with respect to the slow-time index m , such that the channel $h_m(\tau)$ in (2.11) can be written as*

$$h_m(\tau) = \sum_{k=1}^L \alpha_k(mT_{\text{st}}) \delta(\tau - \tau_k(mT_{\text{st}})) + \sum_{k=L+1}^K \alpha_k \delta(\tau - \tau_k). \quad (3.1)$$

Based on Assumption 1, changes of the sampled channel impulse response (CIR) \mathbf{r}_m , defined in (2.8), over slow time m are of particular interest for the detection. Thus, the mean-removed received signal

$$\tilde{\mathbf{R}} = [\tilde{\mathbf{r}}_1 \ \tilde{\mathbf{r}}_2 \ \cdots \ \tilde{\mathbf{r}}_{M-1}] \quad (3.2)$$

is introduced, where $\tilde{\mathbf{r}}_m = \mathbf{r}_m - \frac{1}{M} \sum_{m=0}^{M-1} \mathbf{r}_m$ is the received signal vector \mathbf{r}_m at slow time m after subtracting the mean over m . Let $\tilde{\mathbf{r}} = \text{vec}(\tilde{\mathbf{R}})$ be the equivalent vector representation of $\tilde{\mathbf{R}}$, analogous to \mathbf{r} and \mathbf{R} . Furthermore, let

$$\mathbf{W} = [\mathbf{w}_0 \ \mathbf{w}_1 \ \cdots \ \mathbf{w}_{M-1}] \quad (3.3)$$

denote a matrix of samples the additive noise process $w(t)$ of the noise at the receiver (see Section 2.1) sampled at (absolute) time $t = mT_{\text{st}} + nT_{\text{ft}}$, where $\mathbf{w}_m = [w(mT_{\text{st}}) \ w(mT_{\text{st}} + T_{\text{ft}}) \ \cdots \ w(mT_{\text{st}} + (N-1)T_{\text{ft}})]$ is a vector of all the noise samples obtained at slow time m over fast time $\tau = nT_{\text{ft}}$, $n \in \{0, 1, \dots, N-1\}$. Let $\mathbf{w} = \text{vec}(\mathbf{W})$ be the equivalent vector representation of \mathbf{W} .

To validate Assumption 1, a measurement of a radio channel inside a vacant parked car was performed with a person walking around in the vicinity of the car. Figure 3.1 shows the received signal \mathbf{R} using a raised cosine pulse $s(\tau)$ with 0.5 GHz bandwidth centered at $f_c = 6.5$ GHz with a rolloff factor of $\beta = 0.5$ to sound the channel. The signal was sampled every $T_{\text{ft}} = 1.33$ ns in fast time and $T_{\text{st}} = 0.1$ s in slow time. Some variation can be observed in Figure 3.1b with amplitudes approximately 10 dB above the noise floor. Thus, the model is “wrong”. However,




Figure 3.1: Subfigure (a) shows the received signal \mathbf{R} and (b) the mean-removed received signal $\tilde{\mathbf{R}}$ of an empty car with a person walking around outside close to the car. After subtraction of the mean, some movement can be seen starting from a delay of $\tau \approx 10$ ns. The amplitude of these movements is approximately 40 dB smaller than the amplitude of the static multipath components seen in (a).

“all models are wrong” [1] and the question of interest is whether Assumption 1 leads to a “useful” model. Since the amplitude of these variations is 40 dB less than that of the static components, a dynamic range of at least 40 dB is required from the radar system to resolve these variations. It is challenging to achieve such a high dynamic range with the small UWB nodes typically placed inside a car. Therefore, we conclude that Assumption 1 results in a model useful for cases without such a high dynamic range, i.e. for most cases of practical interest.

3.1.2 Target Model

The development of sophisticated detection algorithms requires a model of how the parameters $\alpha_k(mT_{\text{st}})$ and $\tau_k(mT_{\text{st}})$ of the MPCs interacting with the target, i.e. for $1 \leq k \leq L$, change over slow time m . Expressing the dependence of $\alpha_k(mT_{\text{st}})$ and $\tau_k(mT_{\text{st}})$ for large unpredictable movements of the target, such as limb movements or a repositioning of the torso is intractable. Generally, these large movements will generate comparatively large changes in the CIR that are easy to detect whereas a (for the most part) motionless target, e.g. a sleeping infant, is much harder to detect. To obtain a useful model, we focus on a model for targets with very small movements, e.g. the respiratory chest motion of a person. See Section 3.2.1 or [68] for more details about the received signal obtained in case of different target movements.

A crudely simplified but very useful model for small movements is obtained by modeling the target as a single moving scatter point, that re-emits the signal received at this point with amplitude α_p corresponding to the radar cross section of the person [A]. This model is illustrated in Figure 3.2 and referred to as “pinhole” or “keyhole” channel model [69]. The signal arriving at the point target can be modeled using a multipath propagation model (2.11) to model the “forward” channel $h_f(\tau)$. Similarly, the propagation of the reflected signal to the receive antenna is modeled by another multipath channel denoted as the “backward” channel $h_b(\tau)$. The target channel $h_{p,m}(\tau)$ models the reflection at the target and the clutter channel $h_c(\tau)$ models all propagation paths that do not interact with the target.

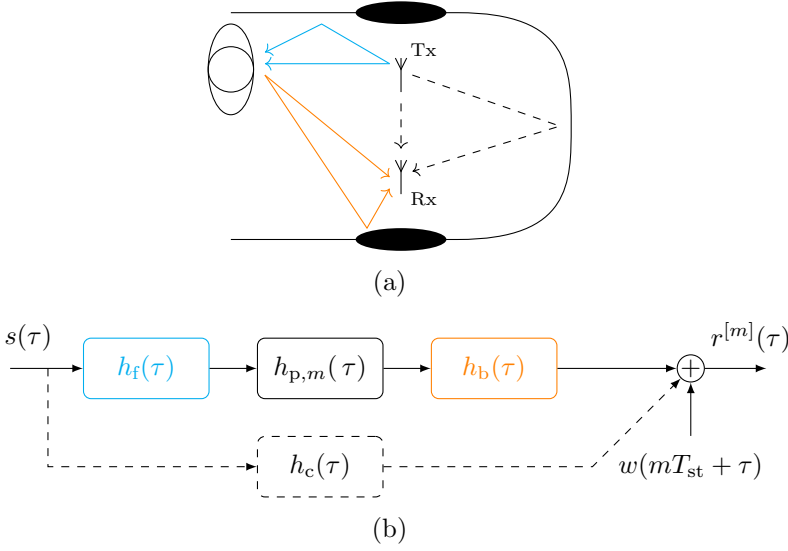


Figure 3.2: Schematic illustration (a) and block diagram (b) of the pinhole channel model for a bistatic setup. For small movements of the target, the only channel that changes over slow time is $h_{p,m}(\tau)$, which models the reflection at the target. The channels $h_f(\tau)$, $h_b(\tau)$ and $h_c(\tau)$ that model the propagation from the transmit antenna to the target, from the target to the receive antenna, and the contribution of all MPCs that do not interact with the target, respectively, are assumed constant with respect to slow time indexed by m .

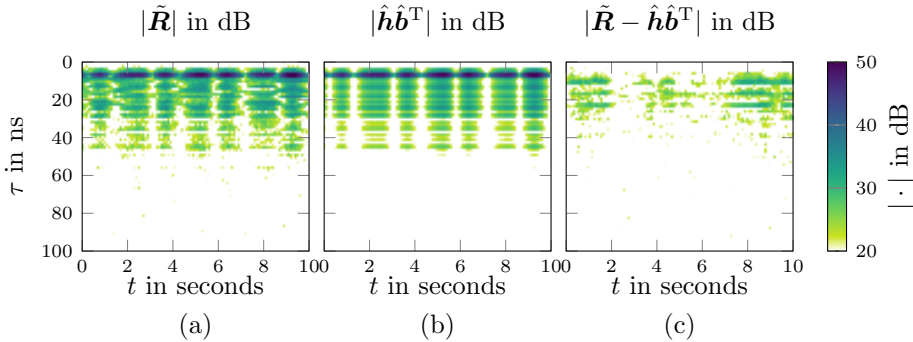


Figure 3.3: Measurement of a person sitting inside the car with only respiratory movement. Subfigure (a) shows the mean removed signal $|\tilde{\mathbf{R}}|$, (b) shows the rank-1 approximation $|\hat{\mathbf{h}}\hat{\mathbf{b}}^T|$ with $\hat{\mathbf{h}}$ and $\hat{\mathbf{b}}$ obtained using an SVD, and (c) shows the difference $|\tilde{\mathbf{R}} - \hat{\mathbf{h}}\hat{\mathbf{b}}^T|$ between the received signal $\tilde{\mathbf{R}}$ and the rank-1 approximation.

Box 3.1: Monostatic, Bistatic and Multistatic Radar.

“Monostatic” refers to a radar system with co-located received and transmit antennas whereas “bistatic” refers to a radar system with the receive and transmit antenna placed at different positions. “Multistatic” refers to radar systems with multiple transmit and/or receive antennas, some but not all of which might be co-located.

Monostatic radar (see Box 3.1) is considered in the following. Nevertheless, the approximation of Assumption 2 can be applied to the bistatic case with minor adjustments. Let $b(t)$ be the movement in direction of the antennas of the scatter point modeling the target relative to its stationary position, $\lambda_c = \frac{f_c}{c}$ be the carrier wavelength and c the propagation speed of the radio waves, i.e. the speed of light. For small movements $|b(t)| \ll \lambda_c$, the (baseband) channels $h_f(\tau)$ and $h_b(\tau)$ can be assumed constant over slow time t except for a phase shift of the carrier wave due to the additional propagation delay $\frac{2b(t)}{c}$ induced by the movement $b(t)$ of the scatter point. This phase shift is modeled as $h_{p,m}(\tau) = \alpha_p \delta\left(\tau - \frac{2b(mT_{st})}{c}\right)$. Inserting $h_{p,m}(\tau)$ into (2.7) results in a multiplication of the pulse $s(\tau)$ with $\alpha_p e^{-j4\pi \frac{b(mT_{st})}{\lambda_c}}$. Based on the pinhole channel model and a small-angle approximation, we derive in [A] that the mean-removed signal can be approximated as a rank-1 matrix in noise as follows.

Assumption 2. *For small movements of the target relative to the carrier wavelength λ_c , e.g. respiratory chest movements, the pinhole channel model illustrated in Figure 3.2 results in the mean-removed signal*

$$\tilde{\mathbf{R}} = \mathbf{h}\mathbf{b}^T + \mathbf{W} \quad (3.4)$$

where $\mathbf{b} = [b(0) \ b(mT_{st}) \ \cdots \ b((M-1)T_{st})]^T \in \mathbb{R}^M$ is a vector of the movement of the target in direction of the antenna $b(t)$ sampled at $t = mT_{st}$, $m = \{0, 1, \dots, M-1\}$, and $\mathbf{h} \in \mathbb{C}^N$ models the combined effects of the forward and backward propagation channels $h_f(\tau)$ and $h_b(\tau)$ up to some multiplicative constant, sampled at $\tau = nT_{ft}$, $n \in \{0, 1, \dots, N-1\}$.

To validate Assumption 2, a measurement was performed with the same setup as the one used to validate Assumption 1, except that an adult was sitting inside the car. The person was motionless, except for the natural respiratory motion of the chest. Figure 3.3 shows the received signal, the rank-1 approximation $\hat{\mathbf{h}}\hat{\mathbf{b}}^T$ obtained using a singular value decomposition (SVD) of $\tilde{\mathbf{R}}$, and the approximation error $\tilde{\mathbf{R}} - \hat{\mathbf{h}}\hat{\mathbf{b}}^T$. Note the channel \mathbf{h} dispersing signal backscattered from the target over fast-time τ . Since Figure 3.3c shows some remaining signal components which cannot be attributed to noise, Assumption 2 does not perfectly approximate the received signal. However, the simple mathematical structure of (3.4) is useful for developing model-based detection algorithms, while, in the shown example, it is able to capture the majority of the signal energy of $\tilde{\mathbf{R}}$ (approximately 93 %).

3.1.3 Energy Detector (ED)

If the car is empty, i.e. if Model \mathcal{H}_0 is true, then no multipath components interact with the target, i.e. $L = 0$ in (3.1). Inserting (3.1) with $L = 0$ via (2.7) and (2.8)

into (2.9), it follows that the mean removed signal $\tilde{\mathbf{r}} = \mathbf{w}$ contains only noise \mathbf{w} , since all terms in \mathbf{r} corresponding to MPCs are constant with respect to the slow-time index m . Assuming \mathbf{w} is AWGN, the mean removed received signal $\tilde{\mathbf{r}}$ is distributed with a circularly-symmetric proper complex Gaussian PDF

$$p(\tilde{\mathbf{r}}|\mathcal{H}_0) = \text{CN}(\tilde{\mathbf{r}}; \mathbf{0}, \lambda^{-1}\mathbf{I}) \propto e^{-\lambda\|\tilde{\mathbf{r}}\|^2} \quad (3.5)$$

where $\lambda \in \mathbb{R}_{\geq 0}$ is the precision (i.e. inverse variance) of the noise samples \mathbf{w} .

A simple approach to car occupancy detection is to calculate the likelihood of $\tilde{\mathbf{r}}$ being generated by $p(\tilde{\mathbf{r}}|\mathcal{H}_0)$. If the noise precision λ is assumed known, then $p(\tilde{\mathbf{r}}|\mathcal{H}_0)$ can be evaluated directly using (3.5). In case the noise precision λ is unknown, a Gamma prior $p(\lambda) = \text{Ga}(\lambda; \epsilon, \eta)$ with shape ϵ and rate η is assumed to obtain $p(\tilde{\mathbf{r}}|\mathcal{H}_0)$ by marginalizing out the unknown noise precision λ using

$$\begin{aligned} p(\tilde{\mathbf{r}}|\mathcal{H}_0) &= \int_0^\infty p(\tilde{\mathbf{r}}|\lambda, \mathcal{H}_0)p(\lambda) d\lambda \\ &\propto (\|\tilde{\mathbf{r}}\|^2 + \eta)^{-(NM+\epsilon)}. \end{aligned} \quad (3.6)$$

The detector

$$T_{\text{ED}}(\tilde{\mathbf{r}}) = -\ln p(\tilde{\mathbf{r}}|\mathcal{H}_0) \underset{\mathcal{H}_1}{\overset{\mathcal{H}_0}{\leq}} \chi_{\text{ED}} \quad (3.7)$$

decides for \mathcal{H}_0 if the (negative) log-evidence $\ln p(\tilde{\mathbf{r}}|\mathcal{H}_0)$ is smaller than some threshold χ_{ED} using either (3.5) or (3.6) for $p(\tilde{\mathbf{r}}|\mathcal{H}_0)$. The negative sign is for consistency with the detectors developed in the following sections, such that a large value of the decision statistics indicates the presence of a target. From (3.5) and (3.6) it can be seen that the decision statistics $T_{\text{ED}}(\tilde{\mathbf{r}})$ of this detector depends only on the energy $\|\tilde{\mathbf{r}}\|^2$ of the (mean-removed) received signal $\tilde{\mathbf{r}}$. Therefore, this detector is referred to as energy detector (ED) [70]. The ED is used as baseline comparison throughout this chapter. Unless otherwise noted, the noise variance is assumed unknown with prior distribution $p(\lambda) \propto \lambda^{-1}$, i.e. Jeffrey's prior, obtained as $(\epsilon, \eta) \rightarrow (0, 0)$.

Assuming all MPC interacting with the target fall within a certain window of delays, i.e. $\tau_{\min} \leq \tau_k(t) \leq \tau_{\max}$ for all $k \leq L$, we can consider only the samples of $\tilde{\mathbf{r}}$ within this window, resulting in the windowed ED [6].

3.1.4 FFT-Detector

A second baseline comparison approach can be obtained from Assumption 2 by further assuming that the channel \mathbf{h} is dominated by a line-of-sight component and the respiratory motion \mathbf{b} is sinusoidal. An intuitive approach to detect such a signal is to perform a fast Fourier transform (FFT) over slow time and compare the largest amplitude against a threshold χ_{FFT} . The FFT over slow time is represented mathematically by right-multiplying the received signal $\tilde{\mathbf{R}}$ with an $M \times M$ discrete Fourier transform (DFT) matrix \mathbf{V}^T with the element at the k th row and n th column of \mathbf{V} given as $[\mathbf{V}]_{k,n} = e^{-j2\pi \frac{(n-1)(k-1)}{M}}$, resulting in the decision statistic

$$T_{\text{FFT}}(\tilde{\mathbf{R}}) = \max_{k,n} |[\tilde{\mathbf{R}}\mathbf{V}^T]_{k,n}| \underset{\mathcal{H}_1}{\overset{\mathcal{H}_0}{\leq}} \chi_{\text{FFT}} \quad (3.8)$$

which is the generalized likelihood ratio test detector, assuming the amplitude, phase and frequency of the sinusoidal breathing motion \mathbf{b} together with the delay of the line-of-sight component of the channel \mathbf{h} are unknown [70].

The assumptions that the breathing motion \mathbf{b} is sinusoidal is quite restrictive and will generally not hold. Similarly, in Figure 3.3a it can be seen that the channel \mathbf{h} contains also MPCs that follow after the line-of-sight component, i.e. the channel \mathbf{h} disperses the target signal. Nevertheless, this model is still useful since the strong assumptions about the shape of the received signal also result in a good detection performance in case those assumptions are (approximately) fulfilled.

3.1.5 Estimator-Correlator (EC)

The ED developed in Section 3.1.3 is very general in the sense that it does not make any assumptions about the received signal in case the target is present, i.e. it is based solely on Assumption 1. On the other hand, the FFT detector makes quite restrictive assumptions about the channel \mathbf{h} and respiratory motion \mathbf{b} . The aim of the estimator-correlator [A] is to exploit the simple mathematical structure of Assumption 2 while being less restrictive in its assumptions about \mathbf{b} and \mathbf{h} as the FFT detector.

From Assumption 2 and an AWGN assumption for the noise \mathbf{w} , it follows that that likelihood of $\tilde{\mathbf{r}}$ under \mathcal{H}_1 is

$$p(\tilde{\mathbf{r}}|\mathbf{b}, \mathbf{h}, \lambda, \mathcal{H}_1) = \text{CN}(\tilde{\mathbf{r}}; \mathbf{b} \otimes \mathbf{h}, \lambda^{-1} \mathbf{I}) \quad (3.9)$$

where \otimes denotes the Kronecker product, and λ the (sample wise) noise precision, i.e. inverse variance, of the circularly-symmetric proper complex Gaussian random variable \mathbf{w} .

The vectors \mathbf{h} and \mathbf{b} are considered unknown for the detection. Assuming independent zero-mean Gaussian prior PDFs $p(\mathbf{b}) = \text{N}(\mathbf{b}; \mathbf{0}, \mathbf{C}_b)$ and $p(\mathbf{h}) = \text{CN}(\mathbf{h}; \mathbf{0}, \mathbf{C}_h)$ with covariances \mathbf{C}_b and \mathbf{C}_h , respectively, the task is to detect a zero-mean signal $\mathbf{x} = \mathbf{b} \otimes \mathbf{h}$ with covariance $\mathbf{C}_x = \mathbf{C}_b \otimes \mathbf{C}_h$ in AWGN. The product $\mathbf{x} = \mathbf{b} \otimes \mathbf{h}$ of two Gaussian random variables does not result in a Gaussian PDF. Nevertheless, a useful model is obtained by assuming that $\mathbf{x} = \mathbf{b} \otimes \mathbf{h}$ is Gaussian distributed with PDF $p(\mathbf{x}) = \text{CN}(\mathbf{x}; \mathbf{0}, \mathbf{C}_x)$. Under this Gaussian assumption, the optimal decision statistic is obtained by the estimator-correlator (EC) as [70]

$$T_{\text{EC}}(\tilde{\mathbf{r}}) = \tilde{\mathbf{r}}^H \mathbf{C}_x (\mathbf{C}_x + \lambda^{-1} \mathbf{I})^{-1} \tilde{\mathbf{r}} \underset{\mathcal{H}_1}{\overset{\mathcal{H}_0}{\leq}} \chi_{\text{EC}} \quad (3.10)$$

which correlates the maximum likelihood estimate $\hat{\mathbf{x}} = \mathbf{C}_x (\mathbf{C}_x + \lambda^{-1} \mathbf{I})^{-1} \tilde{\mathbf{r}}$ of \mathbf{x} with the received signal $\tilde{\mathbf{r}}$ for the detection.

The covariance \mathbf{C}_b can be obtained e.g. by modeling the breathing motion $b(t)$ as a zero-mean Gaussian process with a power spectral density $S_b(f)$ that is constant between $f_{\min} \leq |f| \leq f_{\max}$ and zero everywhere else [A]. The covariance \mathbf{C}_h can be derived from the power-delay-profile of the combined forward-backward channel. See [69] for a detailed characterization of such channels. However, this approach generally depends on the distance between the target and the antennas, which is an additional unknown that has to be considered in the detection process. Alternatively, the covariance can be modeled using a power-delay-profile that follows shape of a

Gamma-distribution with its parameters tuned such that the decay corresponds to the exponential decay we expect and the maximum corresponds to the expected distance from the target to the antenna [B]. Doing so, the decay of the power-delay-profile of the channel can be used as prior without having to explicitly estimate the distance to the target.

3.1.6 Variational Inference (VI)

According to the Bayesian model comparison approach [29, 70], the (log) “Bayes factor” or “odds ratio”

$$T_B(\tilde{\mathbf{r}}) = \ln \frac{p(\tilde{\mathbf{r}}|\mathcal{H}_1)p(\mathcal{H}_1)}{p(\tilde{\mathbf{r}}|\mathcal{H}_0)p(\mathcal{H}_0)} \underset{\mathcal{H}_1}{\overset{\mathcal{H}_0}{\lesseqgtr}} 0 \quad (3.11)$$

should be used to decide between to competing statistical models $p(\tilde{\mathbf{r}}|\mathcal{H}_0)$ and $p(\tilde{\mathbf{r}}|\mathcal{H}_1)$, where $p(\mathcal{H}_1)$ and $p(\mathcal{H}_0) = 1 - p(\mathcal{H}_1)$ are the prior odds of model \mathcal{H}_1 or \mathcal{H}_0 being true. In this work, the priors $p(\mathcal{H}_1)$ and $p(\mathcal{H}_0)$ are tuned to achieve a specified false alarm rate rather than reflecting prior knowledge about the models. If Assumption 1 is applied together with assuming zero knowledge about the distribution of $\tilde{\mathbf{r}}$ in case the target is present, i.e. $p(\tilde{\mathbf{r}}|\mathcal{H}_1) = \text{const.}$, then (3.11) corresponds to the energy-detector (3.7) with $\chi_{\text{ED}} = \ln \frac{p(\mathcal{H}_0)}{p(\mathcal{H}_1)}$.

To increase the performance of the detector, Assumption 2 is applied to obtain the evidence

$$p(\tilde{\mathbf{r}}|\mathcal{H}_1) = \int p(\mathbf{b}, \mathbf{h}, \lambda, \tilde{\mathbf{r}}|\mathcal{H}_1) d\mathbf{b} d\mathbf{h} d\lambda \quad (3.12)$$

for Model \mathcal{H}_1 by marginalizing out the unknowns \mathbf{b} , \mathbf{h} and λ from the joint distribution

$$p(\mathbf{b}, \mathbf{h}, \lambda, \tilde{\mathbf{r}}|\mathcal{H}_1) = p(\tilde{\mathbf{r}}|\mathbf{b}, \mathbf{h}, \lambda, \mathcal{H}_1)p(\mathbf{b})p(\mathbf{h})p(\lambda) \quad (3.13)$$

where $p(\tilde{\mathbf{r}}|\mathbf{b}, \mathbf{h}, \lambda, \mathcal{H}_1)$ given by (3.9) is the likelihood of $\tilde{\mathbf{r}}$ given \mathbf{b} , \mathbf{h} and λ under Assumption 2, and $p(\mathbf{b})$, $p(\mathbf{h})$ and $p(\lambda)$ are the prior distributions for \mathbf{b} , \mathbf{h} and λ , respectively. Unfortunately, solving the integral of (3.12) is intractable. Hence, mean-field VI (see Section 2.3.2) is applied to obtain a lower-bound $\mathcal{L}(q) \leq \ln p(\tilde{\mathbf{r}}|\mathcal{H}_1)$ on the (log) evidence $\ln p(\tilde{\mathbf{r}}|\mathcal{H}_1)$ by approximating the posterior distribution $p(\mathbf{b}, \mathbf{h}, \lambda|\tilde{\mathbf{r}}, \mathcal{H}_1) \approx q(\mathbf{b}, \mathbf{h}, \lambda)$ with a factorized distribution $q(\mathbf{b}, \mathbf{h}, \lambda) = q_b(\mathbf{b}) \cdot q_h(\mathbf{h}) \cdot q_\lambda(\lambda)$ [B].

The ELBO $\mathcal{L}(q)$ is maximized by minimizing the KL-divergence between q and $p(\mathbf{b}, \mathbf{h}, \lambda|\tilde{\mathbf{r}}, \mathcal{H}_1)$. Once the distribution q that maximizes $\mathcal{L}(q)$ is obtained, the Bayesian decision (3.11) is approximated as

$$T_{\text{VI}}(\tilde{\mathbf{r}}) = \mathcal{L}(q) - \ln p(\tilde{\mathbf{r}}|\mathcal{H}_0) \underset{\mathcal{H}_1}{\overset{\mathcal{H}_0}{\lesseqgtr}} \chi_{\text{VI}} \quad (3.14)$$

where $\mathcal{L}(q)$ is calculate using (2.34), $\chi_{\text{VI}} = \ln \frac{p(\mathcal{H}_0)}{p(\mathcal{H}_1)}$ is a threshold depending on the prior probabilities $p(\mathcal{H}_1)$ and $p(\mathcal{H}_0)$, and $p(\tilde{\mathbf{r}}|\mathcal{H}_0)$ is obtained by marginalizing out the unknown noise precision λ as in (3.6) [B]. Note, that from $\mathcal{L}(q) \leq \ln p(\tilde{\mathbf{r}}|\mathcal{H}_1)$ it follows that the decision statistics $T_{\text{VI}}(\tilde{\mathbf{r}})$ is a lower bound on the Bayesian decision statistic $T_B(\tilde{\mathbf{r}})$, i.e. $T_B(\tilde{\mathbf{r}}) \geq T_{\text{VI}}(\tilde{\mathbf{r}})$.

Using conjugate priors, i.e. zero-mean Gaussian priors $p(\mathbf{b}) = \mathcal{N}(\mathbf{b}; \mathbf{0}, \mathbf{C}_b)$ and $p(\mathbf{h}) = \text{CN}(\mathbf{h}; \mathbf{0}, \mathbf{C}_h)$, as in Section 3.1.5, and a Gamma prior $p(\lambda) = \text{Ga}(\lambda; \epsilon, \eta)$, as

in Section 3.1.3, the factors $q_i \in \{q_b, q_h, q_\lambda\}$ resulting from inserting (3.13) into the variational update equation (2.41), are Gaussian distributions $q_b = \mathcal{N}(\mathbf{b}; \hat{\mathbf{b}}, \hat{\mathbf{C}}_b)$ and $q_h(\mathbf{h}) = \mathcal{CN}(\mathbf{h}; \hat{\mathbf{h}}, \hat{\mathbf{C}}_h)$, and a Gamma distribution $q_\lambda = \text{Ga}(\lambda; NM + \epsilon, \hat{M}_\lambda)$ [B]. The parameters $\hat{\mathbf{b}}, \hat{\mathbf{C}}_b, \hat{\mathbf{h}}, \hat{\mathbf{C}}_h$ and $\hat{M}_\lambda = \frac{NM + \epsilon}{\hat{\lambda}}$ of the approximate posterior q are calculated by iteratively applying the recursive update equations

$$\hat{\mathbf{b}} = 2\hat{\lambda}\hat{\mathbf{C}}_b\text{Re}\{\hat{\mathbf{H}}^H\tilde{\mathbf{r}}\} \quad \hat{\mathbf{C}}_b = (\mathbf{C}_b^{-1} + 2\hat{\lambda}\hat{E}_b\mathbf{I})^{-1} \quad (3.15)$$

$$\hat{\mathbf{h}} = \hat{\lambda}\hat{\mathbf{C}}_h\hat{\mathbf{B}}^T\tilde{\mathbf{r}} \quad \hat{\mathbf{C}}_h = (\mathbf{C}_h^{-1} + \hat{\lambda}\hat{E}_h\mathbf{I})^{-1} \quad (3.16)$$

$$\hat{\lambda} = \frac{NM + \epsilon}{\|\tilde{\mathbf{r}}\|^2 - 2\hat{\mathbf{b}}^T\text{Re}\{\hat{\mathbf{H}}^H\tilde{\mathbf{r}}\} + \hat{E}_b\hat{E}_h + \eta} \quad (3.17)$$

until convergence, where $\hat{\mathbf{B}} = \hat{\mathbf{b}} \otimes \mathbf{I}$ and $\hat{\mathbf{H}} = \mathbf{I} \otimes \hat{\mathbf{h}}$ are matrices such that the Kronecker product $\hat{\mathbf{b}} \otimes \hat{\mathbf{h}} = \hat{\mathbf{H}}\hat{\mathbf{b}} = \hat{\mathbf{B}}\hat{\mathbf{h}}$ can be expressed as matrix-vector multiplication, $\hat{E}_b = \|\hat{\mathbf{b}}\|^2 + \text{tr}(\hat{\mathbf{C}}_b)$ is the average signal energy of q_b , and $\hat{E}_h = \|\hat{\mathbf{h}}\|^2 + \text{tr}(\hat{\mathbf{C}}_h)$ is the average signal energy of q_h [B].

Iterating the update equations (3.15)–(3.17) requires matrix inversions. These computationally costly operations can be simplified by performing an eigenvalue decomposition of $\mathbf{C}_b = \mathbf{U}_b\mathbf{S}_b\mathbf{U}_b^H$ and $\mathbf{C}_h = \mathbf{U}_h\mathbf{S}_h\mathbf{U}_h^H$, such that for the calculation of

$$\hat{\mathbf{C}}_b = (\mathbf{C}_b^{-1} + 2\hat{\lambda}\hat{E}_b\mathbf{I})^{-1} = \mathbf{U}_b(\mathbf{S}_b^{-1} + 2\hat{\lambda}\hat{E}_b\mathbf{I})^{-1}\mathbf{U}_b^H \quad (3.18)$$

and

$$\hat{\mathbf{C}}_h = (\mathbf{C}_h^{-1} + \hat{\lambda}\hat{E}_h\mathbf{I})^{-1} = \mathbf{U}_h(\mathbf{S}_h^{-1} + \hat{\lambda}\hat{E}_h\mathbf{I})^{-1}\mathbf{U}_h^H \quad (3.19)$$

only the inverse of diagonal matrices $\mathbf{S}_b^{-1} + 2\hat{\lambda}\hat{E}_b\mathbf{I}$ and $\mathbf{S}_h^{-1} + \hat{\lambda}\hat{E}_h\mathbf{I}$ is required. Furthermore, using $(\mathbf{B}^T \otimes \mathbf{A})\text{vec}(\mathbf{X}) = \text{vec}(\mathbf{A}\mathbf{X}\mathbf{B})$, the products $\hat{\mathbf{H}}^H\tilde{\mathbf{r}} = \tilde{\mathbf{R}}^T\hat{\mathbf{h}}^*$ and $\hat{\mathbf{B}}^T\tilde{\mathbf{r}} = \tilde{\mathbf{R}}\hat{\mathbf{b}}$ can be calculated in an efficient manner, where $\hat{\mathbf{h}}^*$ denotes the (element-wise) complex conjugate of $\hat{\mathbf{h}}$.

Box 3.2 summarizes the model-based detectors developed throughout this chapter as well as the assumptions those detectors are based on.

Box 3.2: Summary of Model-Based Detectors.

Energy Detector

$$T_{\text{ED}}(\tilde{\mathbf{r}}) = -\ln p(\tilde{\mathbf{r}}|\mathcal{H}_0) \underset{\mathcal{H}_1}{\overset{\mathcal{H}_0}{\leq}} \chi_{\text{ED}}$$

Based solely on Assumption 1, i.e. the ED assumes an uninformative model in case the car is occupied. The energy-detector calculates the (log) probability of the data $\tilde{\mathbf{r}}$ being generated under Model \mathcal{H}_0 .

FFT Detector

$$T_{\text{FFT}}(\tilde{\mathbf{R}}) = \max_{k,n} |[\tilde{\mathbf{R}}\mathbf{V}^T]_{k,n}| \underset{\mathcal{H}_1}{\overset{\mathcal{H}_0}{\leq}} \chi_{\text{FFT}}$$

Based on Assumptions 1&2 as well assuming that \mathbf{b} is a sinusoidal signal and \mathbf{h} consists only of a single MPC. The FFT detector picks the slow-time frequency component with largest amplitude and compares it against a threshold.

Estimator Correlator [A]

$$T_{\text{EC}}(\tilde{\mathbf{r}}) = \tilde{\mathbf{r}}^{\text{H}} \mathbf{C}_{\text{x}} (\mathbf{C}_{\text{x}} + \lambda^{-1} \mathbf{I})^{-1} \tilde{\mathbf{r}} \underset{\mathcal{H}_1}{\overset{\mathcal{H}_0}{\leq}} \chi_{\text{EC}}$$

Based on Assumptions 1&2 as well as assuming that λ is known and, \mathbf{b} and \mathbf{h} are random variables with zero mean and covariance \mathbf{C}_{b} and \mathbf{C}_{h} , respectively. The estimator-correlator (EC) estimates a point estimate of $\hat{\mathbf{x}}$ of $\mathbf{x} = \mathbf{b} \otimes \mathbf{h}$ to perform a generalized likelihood ratio test. It considers only the correlation that results from the Kronecker product $\mathbf{b} \otimes \mathbf{h}$ via the covariance $\mathbf{C}_{\text{x}} = \mathbf{C}_{\text{b}} \otimes \mathbf{C}_{\text{h}}$.

Variational Inference [B]

$$T_{\text{VI}}(\tilde{\mathbf{r}}) = \mathcal{L}(q) - \ln p(\tilde{\mathbf{r}}|\mathcal{H}_0) \underset{\mathcal{H}_1}{\overset{\mathcal{H}_0}{\leq}} \chi_{\text{VI}}$$

Based on Assumptions 1&2 as well as assuming that \mathbf{b} and \mathbf{h} are Gaussian distributed random variables and λ is Gamma distributed. The VI-detector approximates the model evidence $p(\tilde{\mathbf{r}}|\mathcal{H}_1)$ used in the Bayesian odds ratio with the ELBO $\mathcal{L}(q)$ based on a factorized approximation of the posterior distribution $p(\mathbf{b}, \mathbf{h}, \lambda|\tilde{\mathbf{r}}, \mathcal{H}_1) \approx q(\mathbf{b}, \mathbf{h}, \lambda) = q_{\text{b}}(\mathbf{b}) \cdot q_{\text{h}}(\mathbf{h}) \cdot q_{\lambda}(\lambda)$.

3.2 Data-driven Car Occupancy Detection

In Section 3.1, car occupancy detection algorithms are derived for the case of small movements. Although these detectors perform well for this case [A,B], there are several drawbacks that can be improved upon: (i) the approximation in the models, (ii) it is hard to specify prior distributions for physiological processes, e.g. the chest motion due to respiration or while talking, and (iii) only small movements are modeled, but no larger movements such as limb movements. To address these issues, a data-driven detection approach is applied which learns to detect car occupants from data using the ResNet convolutional neural network (CNN) architecture [71]. In the learning process, a model of the received signal is implicitly constructed via the weights of the network.

3.2.1 Dataset Description

To gather data for training and the evaluation of the detection algorithms, a measurement campaign was conducted at Graz University of Technology [68]. A total of 34 participants (6 female, 28 male) from the staff and students of the university volunteered to participate in the campaign. CIRs have been recorded with each participant individually sitting inside a car while performing three different activities for a duration of 2 minutes each. The activities are (i) breathing: sitting still with no intentional movement, (ii) talking: talking continuously for the duration of the measurement, e.g. by reading from a book or a freely improvised monologue, and (iii) moving: looking around for small objects hidden inside the car. A maximum length sequence correlative channel sounder with a nominal frequency range of 3.8–10.2 GHz was used to measure the CIRs between one transmit and two receive antennas at a repetition interval of $T_{\text{st}} = 0.1$ s. The data from each participant was segmented into non-overlapping frames of 10 s length each, resulting

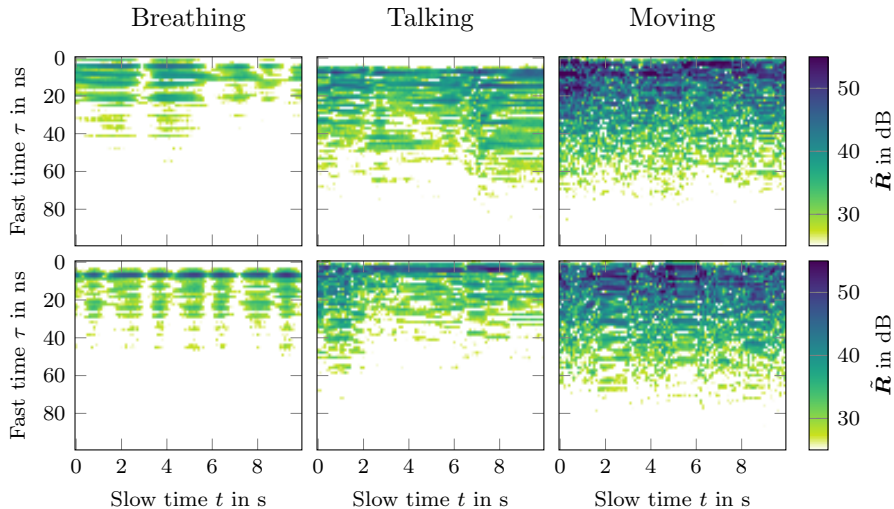


Figure 3.4: Example signals from the dataset. The plots show the magnitude of the received signal \tilde{R} in dB for the breathing activity (left column), talking activity (middle column) and moving activity (right column). The x-axis is the measurement time in seconds (“slow time”) whereas the y-axis is the propagation delay in nanoseconds (“fast time”). The color scale is clipped so that all values ≥ 55 dB map to dark blue (best viewed in color).

in (approximately) 12 samples \mathbf{R} per participant and activity. Data was recorded in two cars, with CIRs from 9 participants being measured in both cars whereas for the remaining participants CIRs have been measured in only one of the cars.

In post processing, each sample \mathbf{R} was multiplied in the frequency domain with a raised cosine pulse with center frequency $f_c = 6.5$ GHz, bandwidth $B = 500$ MHz and roll-off factor $\beta = 0.5$, corresponding to the UWB channel 5 in [4], before being converted down to baseband. Figure 3.4 depicts samples of the dataset. Note, the different energy levels of the signal \tilde{R} which reaches amplitudes of approximately 50 dB for the breathing case, 55 dB for the talking case, and 65 dB for the moving case. Furthermore, a clear structure can be seen in the breathing case, whereas the changes in the signal are less predictable in the talking and moving cases.

3.2.2 ResNet-based Occupancy Detection

Architecture

To detect the occupants, two CNN architectures are trained on the dataset [68] in [C]. Both architectures are based on the popular ResNet architecture [71]. The first architecture takes as input the real and imaginary parts of the complex matrix $\tilde{\mathbf{R}}$ as two different channels, i.e. the input is a tensor of size $N \times M \times 2$. The CNN consists of a 2D-convolutional layer with 16 kernels of size 3×3 operating on the first two dimensions of the input, and batch normalization layer. These initial layers are followed by a repeated structure of pairs of 2D-convolutional layers and batch normalization layers, with skip connections running in parallel to every pair

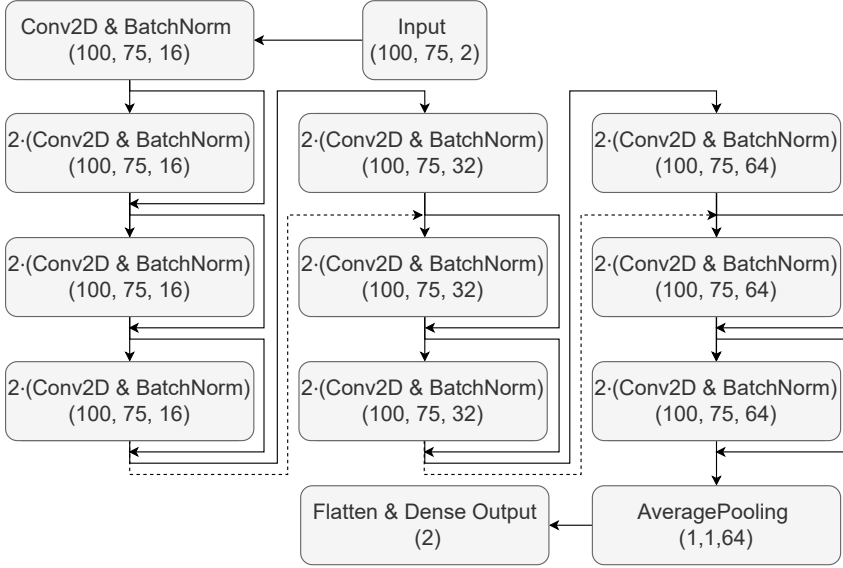


Figure 3.5: Visualization of the 2D-CNN architecture. Connections and skip connections are denoted with arrows and projection skip connections are denoted with dotted arrows. The numbers in brackets denote the output dimensions of a block.

of convolutional layers. We denote such a block of two convolutional layers followed by a batch normalization layer as a “ResNet block”. The number of convolutional kernels (i.e. channels) doubles every 3 ResNet blocks and a total of 9 ResNet blocks is used. The network ends with an average pooling layer across each channel, followed by a flattening and a dense output layer that performs a binary decision between the statistical models \mathcal{H}_0 (the car is empty) and \mathcal{H}_1 (the car is occupied). To preserve the size, we use *same padding*. This 2D-convolutional architecture is shown in Figure 3.5.

As discussed in Section 3.2, the breathing activity results in the signal with the lowest energy, i.e. the signal which is hardest to detect. These signals are characterized by a (quasi) periodic change over slow time. To detect these signals in a computationally efficient manner, the second architecture uses 1D-convolutions with kernel size 3 along slow time while considering the real and imaginary part of each column of $\tilde{\mathbf{R}}$, i.e. different samples of fast time, as $2N$ independent channels, i.e. the input to the network is an $M \times 2N$ matrix obtained by stacking the real and imaginary parts of $\tilde{\mathbf{R}}^T$ along the dimension corresponding to fast time. Besides using 1D-convolutions instead of 2D-convolutions, the second architecture is structured the same as the first architecture. It uses an initial convolution and batch normalization layer to compress the $2N$ input channels to 16 channels followed by 9 ResNet blocks that use 1D-convolutions instead of 2D-convolutions with the number of output channels per ResNet block doubling every 3 ResNet blocks. Finally, the networks ends using an average pooling layer followed by a flatten dense output layer that again performs the final binary decision. This 1D-convolutional architecture is shown in Figure 3.6.

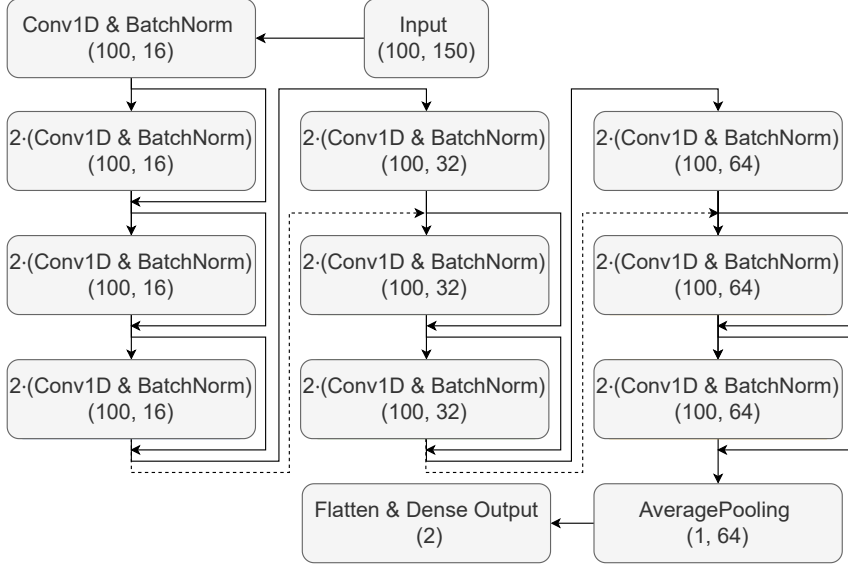


Figure 3.6: Visualization of the 1D-CNN architecture. Connections and skip connections are denoted with arrows and projection skip connections are denoted with dotted arrows. The numbers in brackets denote the output dimensions of a block.

Data Augmentation and Training

The dataset [68] contains data from two different cars and is split into training, validation and test data as follows. For training, only data from the first car are used, while the data from the other car are split into validation and test sets. The last 150 samples of each activity from the second car comprise the test data, while the remaining data from the second car are used as validation set. With this split, data from one car and 13 participants are unseen during training and are only used for the validation and test sets. While the validation set is used to govern early stopping during training, the test set is used only after the architectures are designed and trained.

The signal-to-noise ratio (SNR) of the samples in the dataset is quite high, due to the used laboratory-grade equipment, making the detection trivial. However, the SNR achieved in a real scenario might be smaller, e.g. due to the smaller radar cross section and smaller respiratory movement of infants compared to adults. To train the network to detect occupants also in low-SNR scenarios, the data is augmented by adding white noise with different noise levels to each sample. To preserve the differences in amplitude between the activities, the noise level (and thus SNR) is specified relative to the median energy of the breathing subclass, i.e. $\text{SNR} = \frac{E_b}{2NM\sigma^2}$, where $E_b = \text{median}_{\tilde{\mathbf{R}} \in \mathcal{R}_b}(\|\tilde{\mathbf{R}}\|)$ and \mathcal{R}_b refers to all samples of the breathing activity in the dataset and σ^2 is the sample variance of the real and imaginary parts of the complex AWGN.




Figure 3.7: AUC of the model-based detectors of Section 3.1, summarized in Box 3.2, and the data-driven detector of Section 3.2 as a function of the SNR for the breathing activity of the UWBCarGraz dataset [68].

3.3 Results

The different model-driven detectors of Section 3.1 and the data-driven approach of Section 3.2 are compared using the UWBCarGraz dataset [68]. The area under receiver operating curve (AUC) of the detectors is used as performance metric, because it provides a threshold-independent measure of the performance of a binary classifier. An AUC = 0.5 means the detector output is an arbitrary guess, whereas an AUC = 1.0 indicates that the data is perfectly separable, i.e. that there exists a threshold which achieves perfect detection (i.e. zero misses) while at the same time producing no false alarms.

Figure 3.7 shows the detection performance for the breathing activity for a range of SNRs for all detectors. The performance of the model-driven detectors is evaluated on the whole dataset whereas the performance of the ResNet detector is evaluated only on the test set. The best AUC is achieved by the data-driven ResNet detector [C] where as the best model-driven detector is the VI-detector [B], which outperforms the EC [A] which again outperforms the baseline approaches of the FFT-detector and ED.

Figure 3.8 compares the performance (in terms of AUC over SNR) of the data-driven ResNet-detector [C] against the best model-driven detector, i.e. the VI detector [B]. In this evaluation, both methods are evaluated on the test set only. The evaluation is performed for the three activities separately due to the inherent differences in the data. Due to the larger amplitudes, the talking and moving activities can be detected also in a higher noise level, i.e. at lower SNR. The ResNet detector outperforms the VI detector also in case the target is moving or talking. This is not surprising, given that the Assumption 2 is violated in those activities due to the larger movements of the target.




Figure 3.8: AUC of the ResNet detector [C] and VI detector [B] for the different activities in the UWBCarGraz dataset [68].

3.4 Concluding Remarks

This chapter investigates Hypothesis 1, i.e. the hypothesis that the performance of car occupancy detection algorithms can be improved by taking the multipath channel into account. The results presented Figure 3.7 show, that the model-driven detectors developed in [A,B] which explicitly consider the multipath channel achieve a better detection accuracy in terms of AUC in low-SNR conditions compared to the ED, which considers the multipath channel but makes no assumptions about the target, and the FFT-detector, which assumes the channel consists only of a line-of-sight component but not other MPCs.

The model-driven detectors are based on a signal model for the received signal in case the car is empty/occupied that assumes (i) a static environment and (ii) a point target with small movement. However, since the model is only valid for small movements, a dataset [68] was gathered to train a data-driven detection algorithm based on the ResNet CNN architecture [71]. The ResNet-detector outperforms the model-driven detectors both for large and small target movements, even though the VI-detector developed in Section 3.1.6 approximates the Bayesian-optimal decision. Thus, an interesting topic to explore in further research is investigating why the ResNet-detector is able to outperform the supposedly optimal VI-detector. Possible explanations include (i) model mismatch, e.g. due to the assumptions and approximations made in deriving the model, (ii) the approximation-gap of approximating the true posterior with a variational mean-field approximation and (iii) that the ResNet (implicitly) learns a more specific prior for the channel and respiratory motion from the data.

One interpretation of the results of the model-based detectors is, that the ED is based on a model not specific enough, whereas the FFT-detector is based on a model that is too specific, i.e. based on unrealistic assumptions such as the respiratory movement being (perfectly) sinusoidal and the channel consisting solely of a line-of-sight component. The EC [A] and VI-detector [B] achieve a better trade-off between assumptions which are too specific and unrealistic, and assumptions which

are not specific enough, resulting in low performance. In a Bayesian view, the model used for the FFT-detector can be interpreted as a sort of prior that is concentrated on a small volume but does not reflect the true distribution of the signals well compared to the model of the ED which can be interpreted as a wide prior that covers a much larger volume than the true distribution of the data. Due to the superior performance compared to the ED and FFT-detector, the pinhole channel model used in [A,B] is considered useful, despite the approximations and (partly unrealistic) assumptions it is based on.

Before deploying these detection algorithms to future or existing cars equipped with UWB radios, it should be verified that they generalize well to new scenarios, e.g. to infants, different car models or multiple people occupying the car at the same time. This is especially important for the data-driven ResNet-detector, since it is a “black box” approach driven solely by data from adult participants and only two different car models.

4

Estimation of Structured Line Spectra

In this thesis, a “structured” line spectrum refers to a line spectrum which consists of clusters of related spectral lines. Specifically, we are concerned with estimating the parameters K and $\boldsymbol{\theta} = [\theta_1 \ \theta_2 \ \cdots \ \theta_K]^T$ of the generative model

$$y(t) = \sum_{k=1}^K \sum_{l \in \mathcal{S}_k} \alpha_{k,l} \cdot e^{j2\pi g(\theta_k, l)t} + w(t) \quad (4.1)$$

where $y(t)$ is some signal generated as superposition of K sources, each of which produces a set of correlated lines index by $\mathcal{S}_k \subset \mathbb{Z}$ in AWGN $w(t)$. Each spectral line is weighted with amplitude $\alpha_{k,l} \in \mathbb{C}$, which are of secondary interest. The function $g(\theta_k, l)$ relates the frequency of the l th spectral line of the k th source to a joint parameter θ_k shared between all spectral lines of that source. As shown in Section 4.1, the signal backscattered from extended objects, which disperse the backscattered signal over time, can be approximated with a structured line spectrum. Structured line spectra arise naturally also in other applications, e.g. in multi-pitch estimation of audio data, or estimating the direction of arrival (DOA) of narrowband signals impinging on an antenna array using multiple measurement vectors (MMVs).

Comparing (4.1) to the canonical LSE model (2.12), the main difference is that (2.12) expresses a line spectrum of K lines using K independent parameters for their frequencies, whereas the structured LSE model (4.1) uses K parameters to express a line spectrum consisting of $\sum_{k=1}^K |\mathcal{S}_k| \geq K$ spectral lines. Developing inference algorithms for the structured LSE model incurs additional challenges compared to the canonical LSE model, since each parameter θ_k potentially depends on multiple spectral lines. Furthermore, if the sets \mathcal{S}_k are considered unknown nuisance parameters, we face a non-trivial nested model order estimation problem of finding the sets of spectral lines belonging to each source $\mathcal{S}_k \ \forall k \in \{1, 2, \dots, K\}$ together with the number of sources K .

Sparse Generative Model

Analogous to the case of unstructured LSE, let K_{\max} be a design parameter large enough that $K_{\max} \geq K$. The number of sources is artificially increased to K_{\max} , such that the samples $\mathbf{y} = [y(t_1) \ y(t_2) \ \cdots \ y(t_N)]^T$ obtained from the generative model (4.1) at times $\mathbf{t} = [t_1 \ t_2 \ \cdots \ t_N]^T$ can be expressed as a vector-matrix multiplication

$$\mathbf{y} = \Psi(\boldsymbol{\theta})\boldsymbol{\alpha} + \mathbf{w} \quad (4.2)$$

where $\boldsymbol{\theta} = [\theta_1 \ \theta_2 \ \cdots \ \theta_{K_{\max}}]^T$ is a vector of all group parameters θ_k , $\Psi(\boldsymbol{\theta}) = [\Psi_1(\theta_1) \ \Psi_2(\theta_2) \ \cdots \ \Psi_{K_{\max}}(\theta_{K_{\max}})]$ is a matrix where each block $\Psi_k(\theta_k)$ is composed of the column vectors corresponding to individual spectral lines of the k th source $\Psi_k(\theta_k) = [e^{j2\pi g(\theta_k, l)t} : l \in \mathcal{S}_k]$, $\boldsymbol{\alpha} = [\alpha_1^T \ \alpha_2^T \ \cdots \ \alpha_{K_{\max}}^T]^T$ is the joint vector of all amplitudes, $\boldsymbol{\alpha}_k = [\alpha_{k,l} : l \in \mathcal{S}_k]^T$ is the vector of amplitudes $\alpha_{k,l}$ corresponding to the k th source, and $\mathbf{w} = [w(t_1) \ w(t_2) \ \cdots \ w(t_N)]^T$ are samples of the noise process $w(t)$.

The vector $\boldsymbol{\alpha}$ in (4.2) consists of K nonzero subvectors $\boldsymbol{\alpha}_k$, $1 \leq k \leq K$ and $K_{\max} - K$ subvectors $\boldsymbol{\alpha}_k$, $K < k \leq K_{\max}$ of all zeros. Without loss of generality, we can reorder the elements of $\boldsymbol{\alpha}$ by applying the same reordering also to the columns of $\Psi(\boldsymbol{\theta})$. In case the index set \mathcal{S}_k is unknown, as would be typically the case for the detection of extended objects, we write (4.1) as a sparse matrix/vector product of the form (4.2) by introducing a second design parameter, the maximum viable index set $\mathcal{S}_{\max} \subset \mathbb{Z}$, chosen large enough such that for all $k \in \{1, 2, \dots, K_{\max}\}$: $\mathcal{S}_k \subseteq \mathcal{S}_{\max}$. Using $\alpha_{k,l} = 0$ for $l \in \mathcal{S}_{\max} \setminus \mathcal{S}_k$, we artificially increase the set of spectral lines of each source from \mathcal{S}_k to \mathcal{S}_{\max} by replacing \mathcal{S}_k in the definition of $\Psi_k(\theta_k)$ and $\boldsymbol{\alpha}_k$ with \mathcal{S}_{\max} . The resulting amplitude vector $\boldsymbol{\alpha} \in \mathbb{C}^{|\mathcal{S}_{\max}|K_{\max}}$ has a fixed length. In this case, $\boldsymbol{\alpha}$ is sparse on two levels. First, $\boldsymbol{\alpha}$ consists of K nonzero blocks $\boldsymbol{\alpha}_k \neq \mathbf{0}$ followed by $K_{\max} - K$ blocks of all zeros, i.e. $\boldsymbol{\alpha}_k = \mathbf{0}$ for $k > K$. Second, each of the nonzero blocks $\boldsymbol{\alpha}_k$, $k \leq K$, is itself a sparse vector consisting of $|\mathcal{S}_k|$ nonzero entries and $|\mathcal{S}_{\max}| - |\mathcal{S}_k|$ zeros.

4.1 Modeling Extended Objects Using Structured Line Spectra

Extended objects are defined in [16] as objects occupying more than a single sensor cell. In radar systems applying super-resolution, this definition can not be applied straightforwardly, since there is no discrete grid that is evaluated by the detection algorithm. Therefore, strictly speaking, everything could be considered an extended object. In this thesis, extended objects are defined as objects which can not be modeled sufficiently by a single point-like scattering source or, equivalently, by a single MPC. Note, that the backscatter signal is dispersed over time, since it consists of multiple copies of the transmit pulse $s(t)$, arriving with different delays each.

To derive how extended objects can be modeled using (4.1), consider the case illustrated in Figure 4.1a of a bistatic radar observing a scene with an extended object indexed by k occupying the volume \mathcal{V}_k . An approximate model is obtained by assuming that the transmitted signal is scattered at a (possibly large) number of points distributed within the object volume. This model becomes exact in the limit

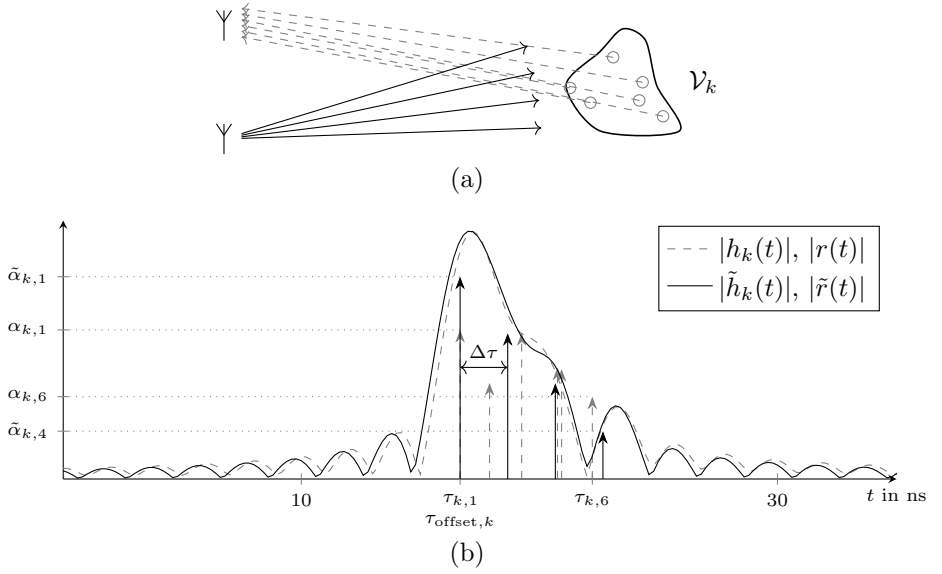


Figure 4.1: Subfigure (a) shows a scene with a single extended object. Subfigure (b) shows the object channel $h_k(t)$, the approximation $\tilde{h}_k(t)$ and the signals $r(t)$ and $\tilde{r}(t)$ obtained by convolution of the pulse $s(t)$ with, respectively, the channel $h_k(t)$ and approximate channel $\tilde{h}_k(t)$. The true channel $h_k(t)$ and signal $r(t)$ are plotted in gray dashed lines. The approximating signal $\tilde{r}(t)$ obtained by approximating $h_k(t)$ with a substitute channel $\tilde{h}_k(t)$ consisting of 4 MPCs equally spaced with distance $\Delta\tau$ is plotted in black.

of considering infinitely many scatter points. Applying the multipath channel model (2.11), each of these scatter points indexed by l contributes an MPC to the channel $h_k(t)$ with amplitude $\alpha_{k,l}$ and delay $\tau_{k,l}$. The received signal $r(t)$ is expressed by (2.4) as the convolution of the transmit pulse $s(t)$ with this channel $h_k(t)$ in noise. In contrast to the previous chapters, the channel is considered stationary and the subscript denoting slow-time is replaced by a subscript indexing different objects. Since the pulse $s(t)$ is typically limited to a certain bandwidth B , we are limited in our capability to resolve these MPCs. Therefore, we approximate the channel with

$$\tilde{h}_k(t) = \sum_{l=0}^{L_k-1} \tilde{\alpha}_{k,l} \cdot \delta(t - \tau_{\text{offset},k} - l\Delta\tau) \quad (4.3)$$

i.e. with a channel consisting of a small sum of L_k substitute MPCs equally spaced with $\Delta\tau = \frac{1}{B}$, where B is the bandwidth of the pulse $s(t)$, corresponding to the time resolution of the system. The delay of the first MPC of the substitute channel is shifted to some delay $\tau_{\text{offset},k}$ equal to the smallest delay of the true channel $h_k(t)$. This approximation is illustrated in Figure 4.1b.

If multiple extended objects are in the scene, we can (approximately) model each of those K objects using a multipath channel $\tilde{h}_k(t)$ of the form (4.3) for $k \in \{1, 2, \dots, K\}$. Assuming a linear medium, the overall (approximate) channel $\tilde{h}(t)$ follows as the superposition of all individual channels, i.e. $\tilde{h}(t) = \sum_{k=1}^K \tilde{h}_k(t)$. The received signal is obtained from (2.4) and (4.3) as

$$r(t) = \sum_{k=1}^K \sum_{l=0}^{L_k-1} \alpha_{k,l} \cdot s(t - \tau_{\text{offset},k} - l\Delta\tau) + w(t) \quad (4.4)$$

where the phase term relating to the carrier frequency is absorbed into the amplitudes $\alpha_{k,l}$ without loss of generality. Transforming (4.4) into frequency domain yields

$$R(f) = \sum_{k=1}^K \sum_{l=0}^{L_k-1} \alpha_k \cdot S(f) \cdot e^{-j2\pi(\tau_{\text{offset},k} + l\Delta\tau)f} + W(f) \quad (4.5)$$

revealing the connection to the structured LSE problem (4.1), since (4.5) is an instance of (4.1), where $g(\theta_k, l) = \theta_k + l\Delta\tau$ and $\mathcal{S}_k = \{0, 1, \dots, L_k - 1\}$, except for the multiplication with the (known) spectrum $S(f)$ of the pulse $s(t)$.

4.2 Fast Inference With Known Group Sizes

4.2.1 Introduction and Related Methods

Several methods have been proposed in the literature to find the support of a group-sparse weight vector in a linear regression problem such as (4.2). Most of these methods do not consider the estimation of $\boldsymbol{\theta}$ on a continuum and instead construct the dictionary matrix $\boldsymbol{\Psi}(\boldsymbol{\theta})$ using a fixed grid $\boldsymbol{\theta} \triangleq \hat{\boldsymbol{\theta}}$. Notable instances of such methods include group-LASSO [72–74], block matching pursuit [75], group-SPICE [76], M-SBL [77], and block-SBL (BSBL) [27, 28]. In the following, this thesis focuses on SBL as an instance of Type-II methods, since Type-II methods include a

Box 4.1: Real- and Complex-Valued Signal Models.

Many SBL algorithms considering block-sparsity, e.g. the BSBL algorithm [27, 28] are derived using a real-valued signal model whereas LSE problems typically rely on a complex-valued signal model [46, 47, 54]. In [D] the signal model is parameterized to cover both the real- and complex-valued cases to enable straightforward comparison of both cases. For brevity, this thesis presents the fast variational block-SBL [D] only for the complex-valued signal model. The results for the real-valued signal model (obtained by replacing all complex-Gaussian distributions with their real-valued equivalents) are, broadly speaking, the same with only minute differences.

more general class of priors compared to Type-I methods. The interested reader is referred to [57] for a detailed discussion about Type-I and Type-II methods.

To find the support of α , group sparsity-based SBL methods such as [27, 28, 77–80] introduce a hierarchical prior $p(\alpha|\gamma)$, which depends on hyperparameters $\gamma = [\gamma_1 \ \gamma_2 \ \cdots \ \gamma_{K_{\max}}]^T \in \mathbb{R}_{>0}^{K_{\max}}$ as

$$p(\alpha|\gamma) = \prod_{k=1}^{K_{\max}} p(\alpha_k|\gamma_k) = \prod_{k=1}^{K_{\max}} \text{CN}(\alpha_k; \mathbf{0}, (\gamma_k \mathbf{D}_k)^{-1}) \quad (4.6)$$

where \mathbf{D}_k is a known Hermitian-symmetric matrix characterizing the intra-block correlation (or a real-valued equivalent of (4.6), see Box 4.1). The support of α is obtained together with an estimate $\hat{\theta}$ of the parameters θ by maximizing the marginal likelihood

$$(\hat{\theta}, \hat{\gamma}) = \arg \max_{\theta, \gamma > 0} \ln \int p(\mathbf{y}|\alpha, \theta, \gamma) p(\alpha|\gamma) d\alpha. \quad (4.7)$$

Several methods can be used to maximize the marginal integral (4.7). For example [27, 28, 77] use an EM algorithm while considering a fixed dictionary matrix $\Psi \triangleq \Psi(\theta)$. However, these methods converge slowly due to the slow convergence of the hyperparameters γ . Following the fast-SBL approach [44], equation (4.7) can be maximized by coordinate ascent for one element $\hat{\gamma}_k$ of $\hat{\gamma}$ at a time. References [78, 79] present a method to check for convergence/divergence of each $\hat{\gamma}_k$ in a single step by solving for the roots of a polynomial of order $2|\mathcal{S}_k| - 1$.

Another approach to SBL introduced in Section 2.2.1 is to estimate posterior distributions $q_{\gamma,k}(\gamma_k)$ instead of point estimates $\hat{\gamma}_k$. Such a variational SBL variant considering group-sparsity was developed in [80]. Introducing a prior PDF $p(\gamma)$ on the hyperparameters γ , [80] approximates the posterior distribution $p(\alpha, \gamma|\mathbf{y}) \propto p(\mathbf{y}|\alpha) p(\alpha|\gamma) p(\gamma)$ with a factorized distribution $q(\alpha, \gamma) = q_\alpha(\alpha) \prod_{k=1}^{K_{\max}} q_{\gamma,k}(\gamma_k)$ using a variational mean-field approximation while considering a fixed dictionary matrix $\Psi(\bar{\theta})$. Due to its iterative nature, the variational approach to SBL exhibits slow convergence similar to EM-based SBL variants. For the non-grouped approach, i.e. for $|\mathcal{S}_k| = 1 \ \forall \ k$, a single-step convergence check for variational SBL was developed in [53]. In [D], we present a novel generalization of the fast variational SBL [53] to the group-sparse case.

4.2.2 Fast Variational Block SBL

To keep the notation concise and focus on the group-sparse fast convergence check, we consider in this subsection the case of \mathcal{S}_k known and a fixed dictionary matrix $\Psi \triangleq \Psi(\bar{\theta})$, restricting the values of θ to be elements of some fixed grid $\bar{\theta}$ with K_{\max} elements.

In [D], the posterior $p(\alpha, \gamma | \mathbf{y})$ is approximated with a factorized distribution $q(\alpha, \gamma) = q_\alpha(\alpha) \prod_{k=1}^{K_{\max}} q_{\gamma,k}(\gamma_k)$ by maximizing the ELBO $\mathcal{L}(q)$ using variational mean-field inference as described in Section 2.3.2. The joint distribution of all variables α , γ and \mathbf{y} is modeled as

$$p(\alpha, \gamma, \mathbf{y}) = p(\mathbf{y} | \alpha) p(\alpha | \gamma) p(\gamma). \quad (4.8)$$

We assume a Gaussian likelihood $p(\mathbf{y} | \alpha) = \text{CN}(\mathbf{y}; \Psi\alpha, \lambda^{-1}\mathbf{I})$ with noise precision λ ,¹ together with the group-sparse SBL prior (4.6). As prior for the hyperparameters γ we use $p(\gamma) = \prod_{k=1}^{K_{\max}} \text{GIG}(\gamma_k; a, b, c)$, where

$$\text{GIG}(\gamma; a, b, c) \propto \gamma^{(c-1)} e^{-\frac{1}{2}(a\gamma + b\gamma^{-1})} \quad (4.9)$$

is the PDF of a generalized inverse Gaussian distribution [80, 81] with domain $\gamma \in \mathbb{R}_{>0}$ and parameters $a > 0$, $b > 0$ and $c \in \mathbb{R}$. This choice of hyperprior $p(\gamma_k)$ includes the popular Gamma distribution, inverse Gamma distribution and Jeffrey's prior $p(\gamma_k) \propto \gamma_k^{-1}$ in the limit of its parameters [D, 80]. Inserting $q(\alpha, \gamma)$ and (4.8) into the variational mean-field update equation (2.41), the distribution q_α which maximizes the ELBO is found to be a normal distribution $q_\alpha(\alpha) = \text{CN}(\alpha; \hat{\alpha}, \hat{\Sigma})$ with mean

$$\hat{\alpha} = \lambda \hat{\Sigma} \Psi^H \mathbf{y} \quad (4.10)$$

and covariance

$$\hat{\Sigma} = (\lambda \Psi^H \Psi + \mathbf{D}_{\hat{\gamma}})^{-1} \quad (4.11)$$

where $\mathbf{D}_{\hat{\gamma}}$ is a hermitian-symmetric block-diagonal matrix with the principal diagonal blocks $\hat{\gamma}_k \mathbf{D}_k$ and $\hat{\gamma}_k = \langle \gamma_k \rangle_{q_{\gamma,k}}$ is the mean of the factor $q_{\gamma,k}(\gamma_k)$ of the approximating distribution $q(\alpha, \gamma)$ [D]. Similarly, the factors $q_{\gamma,k}$ of the approximating distribution which maximize the ELBO $\mathcal{L}(q)$ are shown to be generalized inverse Gaussian distributions

$$q_{\gamma,k}(\gamma_k) = \text{GIG}(\gamma_k; \hat{a}_k, b, \hat{c}_k) \quad (4.12)$$

where $\hat{c}_k = |\mathcal{S}_k| + c$ and

$$\hat{a}_k = 2 \langle \alpha_k^H \mathbf{D}_k \alpha_k \rangle_{q_\alpha} + a. \quad (4.13)$$

Since \mathcal{S}_k is considered known, $\hat{c}_k = |\mathcal{S}_k| + c$ is a fixed quantity. Hence, each distribution $q_{\gamma,k}$ depends only on a single parameter \hat{a}_k , which is fully defined by the mean

$$\hat{\gamma}_k = \sqrt{\frac{b}{\hat{a}_k}} \frac{K_{\hat{c}_k+1}(\sqrt{\hat{a}_k b})}{K_{\hat{c}_k}(\sqrt{\hat{a}_k b})} \quad (4.14)$$

¹For conciseness, it is assumed here that the noise precision λ is known, whereas λ is considered unknown in [D] and included in the estimation procedure.

where $K_c(\cdot)$ denotes the modified Bessel function of the second kind. Equation (4.14) can be simplified for certain choices of a , b , and c , e.g. corresponding to the Gamma distribution or Jeffrey's prior [D, 80].

Following the approach of [53], updates of q_{α} and $q_{\gamma,k}$ repeated ad infinitum as $q_{\alpha} \rightarrow q_{\gamma,k} \rightarrow q_{\alpha} \rightarrow q_{\gamma,k} \rightarrow \dots$, are investigated for a single parameter $\hat{\gamma}_k$. Inserting (4.10) and (4.11) via $q_{\alpha}(\alpha) = \text{CN}(\alpha; \hat{\alpha}, \hat{\Sigma})$ into (4.13), we obtain a recurrent relation $\hat{\gamma}_k^{[n+1]} = f_k(\hat{\gamma}_k^{[n]})$ that expresses the mean $\hat{\gamma}_k^{[n+1]}$ at the $(n+1)$ th iteration of this update loop as a function of the estimate $\hat{\gamma}_k^{[n]}$ at the n th iteration. Consequently, any maximum of the ELBO with finite estimate $\hat{\gamma}_k < \infty$ must be a stable fixed point γ^* fulfilling $\gamma^* = f_k(\gamma^*)$. Furthermore, we show in [D] that the recurrent relation f_k is a strictly increasing rational function, for common choices of the parameters a , b and c of the priors $p(\gamma_k)$. E.g., Jeffrey's prior obtained in the limit as $(a, b, c) \rightarrow (0, 0, 0)$ yields

$$f_k(\gamma) = \frac{|\mathcal{S}_k| A_k(\gamma)}{B_k(\gamma)} \quad (4.15)$$

where $A_k(\gamma)$ and $B_k(\gamma)$ are two polynomials in γ . Inserting (4.15) into the fixed point equation $\gamma^* = f_k(\gamma^*)$, we obtain the fixed points γ^* as solutions to the polynomial equation

$$|\mathcal{S}_k| A_k(\gamma) - \gamma B_k(\gamma) = 0 \quad (4.16)$$

i.e. as roots of the polynomial $G_k(\gamma) = |\mathcal{S}_k| A_k(\gamma) - \gamma B_k(\gamma)$.

In general, there might be multiple viable fixed points, i.e. the polynomial $G_k(\gamma)$ might have several roots in $\mathbb{R}_{>0}$. Since f_k is shown to be strictly monotonic, [D, Theorem 1] determines from the initial estimate $\hat{\gamma}_k^{[0]}$ whether the update sequence $(\hat{\gamma}_k^{[0]}, \hat{\gamma}_k^{[1]}, \dots)$ converges, or if it diverges. If the update sequence converges, [D, Theorem 1] also determines to which fixed point the sequence converges. Specifically, if the first iteration is decreasing, i.e. if $\hat{\gamma}_k^{[1]} \leq \hat{\gamma}_k^{[0]}$, then the sequence always converges to the largest fixed point γ^* that is smaller than or equal to the initial estimate $\hat{\gamma}_k^{[0]}$. If the sequence is increasing, i.e. $\hat{\gamma}_k^{[1]} > \hat{\gamma}_k^{[0]}$, then it converges to the smallest fixed point γ^* larger than $\hat{\gamma}_k^{[0]}$. If no such fixed point exists, the sequence diverges. In summary, [D, Theorem 1] generalizes [53, Theorem 1] to the group-sparse case and provides a fast update rule for $\hat{\gamma}_k$.

The variational fast update corresponds to repeating individual updates of $q_{\alpha}(\alpha)$ and $q_{\gamma,k}(\gamma_k)$. Since each of those individual updates is guaranteed to increase the ELBO, the fast update is also guaranteed to increase the ELBO. Thus, the algorithm is guaranteed to converge to a local optimum of the ELBO. Furthermore, we show in [D] that for Jeffrey's prior, the polynomial $G_k(\gamma_k)$ is the same as the polynomial obtained in [78, 79] for the fast convergence check based on maximizing the marginal likelihood (4.7) via coordinate ascent, i.e. that both methods result in equivalent algorithms.

Thresholding and Increased Sparsity

If SBL is used for LSE, it is shown to overestimate the model order, resulting in additional artificial components [60, 82]. References [60, 82] introduce an additional

threshold into the fast update developed by Tipping [44], and derive the relationship between the threshold and false alarm rate. For variational SBL, a similar modification of the fast update rule was proposed in [53]. Any fixed point γ^* fulfilling $\gamma^* = f_k(\gamma^*)$ is stable if, and only if,

$$\left| \frac{df_k(\gamma)}{d\gamma} \right|_{\gamma=\gamma^*} < 1. \quad (4.17)$$

Reference [53] introduces a threshold $\chi \leq 1$ into (4.17), i.e. it requires

$$\left| \frac{df_k(\gamma)}{d\gamma} \right|_{\gamma=\gamma^*} < \chi. \quad (4.18)$$

It is shown in [53] that this heuristic increases the sparsity of the solution and reduces the detection of spurious artificial components similar to the approach of [60, 82]. In [D], we apply the variational thresholding rule (4.18) to the group-sparse case. Note, that choosing $\chi < 1$ might result in an update which decreases the ELBO. Hence, the algorithm is not guaranteed to converge for $\chi < 1$. Nevertheless, in [D] we did not observe any problems with convergence when using $\chi < 1$.

4.2.3 Results

The algorithm presented in [D] is evaluated on a real-valued toy example using a random dictionary matrix with $2N$ columns grouped into $K_{\max} = \frac{2N}{d}$ groups of equal size $d = 10$. We assume $N = 200$ measurements are obtained unless otherwise stated. The elements of the dictionary matrix are drawn from independent zero-mean normal distributions with unit variance and each column of Ψ is normalized to unit norm. For the amplitude vector α , we selected a number of groups α_k at randomly chosen locations such that the desired sparsity ratio $\frac{\|\alpha\|_0}{N} = 0.2$ where $\|\cdot\|_0$ is the ℓ^0 norm, i.e. the number of nonzero elements of α . The noise precision λ is chosen, such that the SNR defined as $\text{SNR} = \frac{\lambda \|\Psi\alpha\|^2}{N}$ equals 15 dB. As performance metric we use the normalized mean squared error (NMSE) defined as $\text{NMSE} = \frac{\|\alpha - \hat{\alpha}\|^2}{\|\alpha\|^2}$, averaged over 100 simulation runs. See [D] for further simulation details.

Two variants of the fast variational block-SBL are evaluated. Once, with Jeffrey's prior obtained by $(a, b, c) \rightarrow (0, 0, 0)$ and a threshold $\chi = 0.67$, and once with the prior $p(\gamma_k) = \text{const.}$ obtained by $(a, b) \rightarrow (0, 0)$ and $c = 1$ denoted as "scaled Jeffrey's prior," since it corresponds to the prior $p(\gamma_k) \propto \gamma_k^{c-1}$ resembling Jeffrey's prior $p(\gamma_k) \propto \gamma_k^{-1}$ multiplied with the scaling factor γ_k^c . For comparison, we evaluate two variants of the BSBL algorithm [28] denoted as BSBL-BO and BSBL-EM. Additionally, we also evaluate the oracle minimum mean squared error (MMSE) estimator, which knows the location of the nonzero groups in advance and obtains the MMSE estimates $\hat{\alpha}$. For all algorithms we consider $D_k = I$ known.

Figures 4.2a, 4.2b and 4.2d depict the normalized mean squared error (NMSE) of the algorithms as a function of the signal length N , the SNR, and block size d , whereas Figure 4.2c shows the runtime of the algorithms as function of the signal length N . For $N = 500$, the proposed method is faster by approximately 2 orders of magnitude compared to the two BSBL variants [28] while achieving an NMSE virtually identical to that of the oracle MMSE estimator. As shown in Figures 4.2a, 4.2b and 4.2d, the NMSE performance of the oracle estimator is matched

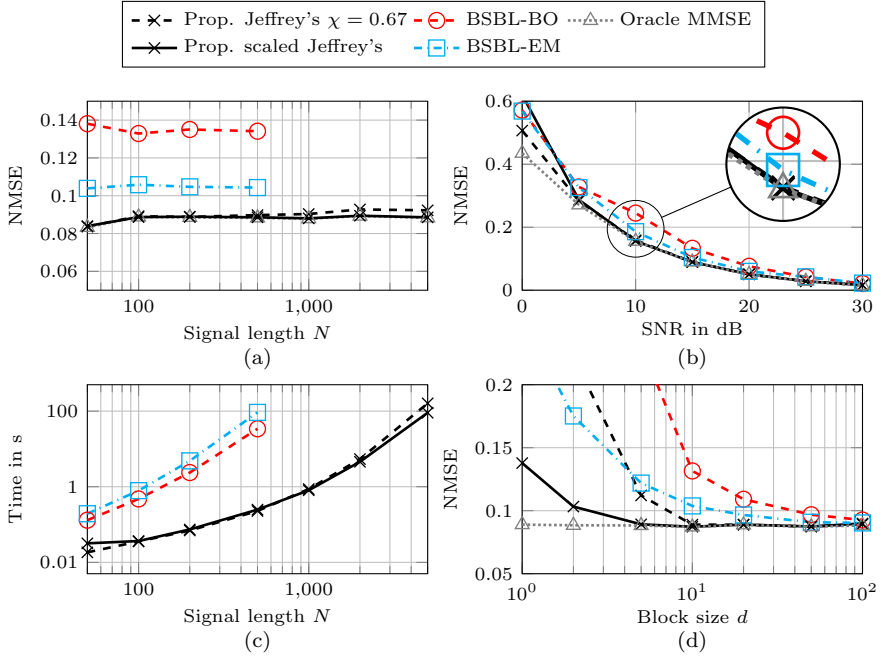


Figure 4.2: Numerical Evaluation of the fast variational block-SBL algorithm [D].

over a wide range of signal lengths N , SNRs, and block sizes d , demonstrating the state-of-the-art performance and superior computational efficiency of the proposed fast variational block-SBL algorithm. We refer the reader to [D] for a more detailed evaluation of the fast variational block-SBL algorithm.

4.2.4 Application: Multiple Measurement Vector (MMV)-based Direction-of-Arrival (DOA) Estimation

One example where structured line spectra with known group size arise naturally is the estimation of the direction of arrivals (DOAs) of signals from multiple narrowband sources impinging on an antenna or microphone array based on multiple measurement vectors (MMVs). Let $\phi(\theta)$ be a column vector containing the array response to a signal impinging from direction θ . Estimating the number of sources and their DOAs based on a single snapshot can be formulated as a canonical LSE problem of the form (2.17) with dictionary $\Phi(\theta) = [\phi(\theta_k) : k \in \{1, 2, \dots, K_{\max}\}]$. To obtain a fixed dictionary, we assume that the DOAs θ_k are elements of some fixed grid $\bar{\theta} = [\bar{\theta}_1 \ \bar{\theta}_2 \ \dots \ \bar{\theta}_{K_{\max}}]^T$ with K_{\max} elements covering the field-of-view of the array, resulting in the dictionary matrix $\Phi \triangleq \Phi(\bar{\theta})$. To consider MMV, let \mathbf{y}_l , $l \in \{1, 2, \dots, L\}$ be multiple successive measurement vectors (i.e. snapshots). Assuming the DOAs are constant over the snapshots, stacking the observations \mathbf{y}_l into a matrix $\mathbf{Y} = [\mathbf{y}_1 \ \mathbf{y}_2 \ \dots \ \mathbf{y}_L]$ yields

$$\mathbf{Y} = \Phi \mathbf{A} + \mathbf{W} \quad (4.19)$$

where $\mathbf{A} = [\boldsymbol{\alpha}_1 \ \boldsymbol{\alpha}_2 \ \cdots \ \boldsymbol{\alpha}_{K_{\max}}]^T$ is a $K_{\max} \times L$ matrix where the k th row consists of the amplitudes $\boldsymbol{\alpha}_k = [\alpha_{k,1} \ \alpha_{k,2} \ \cdots \ \alpha_{k,L}]^T$ of the k th source at snapshot $l \in \{1, 2, \dots, L\}$, and $\mathbf{W} = [\mathbf{w}_1 \ \mathbf{w}_2 \ \cdots \ \mathbf{w}_L]$ is an $N \times L$ matrix obtained by stacking the observation noise \mathbf{w}_l at snapshot l . If an unknown number of sources $K < K_{\max}$ are present, then the matrix \mathbf{A} is row-sparse, i.e. it consists of K nonzero rows and $K_{\max} - K$ rows of zeros. The indexes of the nonzero rows corresponds to the indices of the elements of the grid $\bar{\boldsymbol{\theta}}$ from which signal impinge on the array. The matrix signal model (4.19) can be transformed into a group-sparse vector signal model of the form (4.2) using $\mathbf{y} = \text{vec}(\mathbf{Y}^T)$, $\boldsymbol{\alpha} = \text{vec}(\mathbf{A}^T)$, $\mathbf{w} = \text{vec}(\mathbf{W}^T)$ and $\boldsymbol{\Psi} = \boldsymbol{\Phi} \otimes \mathbf{I}_L$. In this case, the length of each group of amplitudes $\boldsymbol{\alpha}_k$ is equal to the number of snapshots L and, thus, $\mathcal{S}_k = \{1, 2, \dots, L\}$ is known for all k . Hence, the fast variational block SBL [D] can be applied to estimate the number of sources K as well as their DOAs $\boldsymbol{\theta}$ and the amplitudes $\boldsymbol{\alpha}$.

We compare the fast variational block SBL algorithm against the DOA-SBL algorithm [54], since the BSBL variants used as comparison in Section 4.2.3 are derived for a real-valued signal model only. The DOA-SBL algorithm does not perform model order estimation, rather it produces a spectrum with peaks at the locations of potential sources. Hence, two versions of the DOA-SBL algorithm are evaluated, once given the true number of sources K , where the positions of the K largest peaks are taken as estimates $\hat{\boldsymbol{\theta}}$ referred to as “oracle DOA-SBL”, and once by considering each peak larger than some threshold. As performance metric we use the optimal subpattern assignment (OSPA) metric [83], which jointly penalizes errors in estimating the DOAs $\boldsymbol{\theta}$ and errors in estimating the model order. We simulate a uniform linear array with $N = 100$ elements spaced with halve of the carrier wavelength. From this array, $L = 10$ snapshots are obtained with $K = 3$ sources impinging on the array from DOAs $\boldsymbol{\theta} = \{-2^\circ, 3^\circ, 50^\circ\}$. The amplitudes are generated independently between sources. For each source, the amplitudes are modeled over snapshots by a first-order auto-regressive process as $\alpha_{k,l} = \beta \alpha_{k,l-1} + \xi$, where $\alpha_{k,l-1}$ is the amplitude at the previous snapshot and ξ is a zero-mean proper complex Gaussian random variable with independent realizations for each source and snapshots. With this setup, the correlation between the source amplitudes can be varied by varying the parameter β . Three cases are simulated. (i) no correlation using $\beta = 0$, (ii) medium correlation using $\beta = 0.5$ and (iii) strong correlation using $\beta = 0.95$. This intra-block correlation is modeled in the fast variational block SBL using the precision matrix \mathbf{D}_k of the prior $p(\boldsymbol{\alpha}_k | \gamma_k)$.

Figure 4.3 shows the OSPA obtained from the fast variational block SBL algorithm for the three cases together with the performance of the DOA-SBL algorithm as a function of the SNR. The DOA-SBL algorithm can not consider the correlation between the snapshots. Therefore its estimation performance is the same in all three cases. For the case (i) of no correlation ($\beta = 0$), both the fast variational block SBL and the DOA-SBL achieve a similar performance, whereas the oracle DOA-SBL performs best due to the unrealistic assumption of knowing the number of sources. However, in case (ii) and (iii) of $\beta > 0$, the fast variational block SBL achieves a performance gain by considering the correlation. The fast variational block SBL is able to reach a similar OSPA as the oracle DOA-SBL for $\beta = 0.95$. See [D] for further details of how the fast variational block SBL can be applied to MMV DOA estimation. It is expected that the estimation accuracy of the fast variational block SBL improves further if $\boldsymbol{\theta}$ is estimated on a continuum instead of




Figure 4.3: OSPA of the method proposed in [D] and the DOA-SBL algorithm [54] as function of the SNR given $d = 10$ snapshots with different source correlations β .

restricting the estimates to a predefined grid. Estimating θ on a continuum can be achieved using a variational EM algorithm together with fast updates, similar to [47]

4.3 Super-Resolution Inference With Unknown Group Sizes

While Subsection 4.2 considers the case of \mathcal{S}_k known $\forall k$, there are many applications where \mathcal{S}_k is unknown, including the detection and estimation of dispersed backscatter signals from extended objects. Hence, this section details a variational-EM algorithm presented in [E], which estimate the parameters of the structured LSE model (4.1). Specifically, the number of sources K , the source parameters θ , the amplitude vector α , and the index sets $\mathcal{S}_k \forall k$ are estimated. The nested model-order estimation problem of estimating all sets \mathcal{S}_k along with the number of sources K makes the group-sparse LSE problem more challenging compared to the canonical LSE problem or the case for known group sizes considered in Section 4.2. Almost no prior work for structured line spectra can be found, which estimates the parameters θ on a continuum while also estimating K and \mathcal{S}_k . The only other algorithm found in the literature is the block-sparse super-resolution iterative reweight (BSURE-IR) algorithm [84], designed for the problem of multi-pitch estimation.

4.3.1 Variational-EM Approach

Prior Model for Unknown Group Sizes

As generative model, we consider a maximum viable set of spectral lines \mathcal{S}_{\max} , such that $\mathcal{S}_k \subseteq \mathcal{S}_{\max}$ for each source indexed by $k \in \{1, 2, \dots, K_{\max}\}$, to express

(4.1) as a multiplication (4.2) of a dictionary matrix $\Psi(\theta) \in \mathbb{C}^{N \times K_{\max} |S_{\max}|}$ and an amplitude vector $\alpha \in \mathbb{C}^{K_{\max} |S_{\max}|}$ with fixed size each. To perform SBL-like Type-II inference on this model, we extend the hierarchical prior of the super fast SBL algorithm [85] to the group-sparse case. Specifically, we model the prior of the amplitude $\alpha_{k,l}$ of the l th spectral line of the k th source to depend on an independent hyperparameter $\gamma_{k,l} \in \mathbb{R}_{>0}$ and a binary random variable $z_k \in \{0, 1\}$ modeling the existence of the k th group, resulting in the hierarchical prior $p(\alpha_{k,l} | \gamma_{k,l}, z_k) = \text{CN}(\alpha_{k,l}; 0, \gamma_{k,l}^{-1})^{z_k} \delta(\alpha_{k,l})^{1-z_k}$. The joint prior of all amplitudes α is modeled as

$$p(\alpha | \gamma, \mathbf{z}) = \prod_{k=1}^{K_{\max}} \left(\prod_{l \in S_{\max}} \text{CN}(\alpha_{k,l}; 0, \gamma_{k,l}^{-1}) \right)^{z_k} \cdot (\delta(\alpha_k))^{1-z_k} \quad (4.20)$$

where $\mathbf{z} = [z_1 \ z_2 \ \cdots \ z_{K_{\max}}]^T$, $\gamma = [\gamma_1^T \ \gamma_2^T \ \cdots \ \gamma_{K_{\max}}^T]^T$ and $\gamma_k = [\gamma_{k,l} : l \in S_{\max}]^T$. Introducing priors $p(\theta)$, $p(\gamma) = \prod_{k=1}^{K_{\max}} \prod_{l \in S_{\max}} \text{Ga}(\gamma_{k,l}; \eta, \nu)$ and $p(\mathbf{z}) = \prod_{k=1}^{K_{\max}} \bar{z}^{z_k} (1 - \bar{z})^{1-z_k}$ with parameters $\eta \in \mathbb{R}_{>0}$, $\nu \in \mathbb{R}_{>0}$ and $\bar{z} \in [0, 1]$, we model the joint likelihood as

$$p(\alpha, \theta, \gamma, \mathbf{z}, \mathbf{y}) = p(\mathbf{y} | \alpha, \theta) p(\alpha | \gamma, \mathbf{z}) p(\gamma) p(\mathbf{z}) p(\theta) \quad (4.21)$$

where $p(\mathbf{y} | \alpha, \theta) = \text{CN}(\mathbf{y}; \Psi(\theta)\alpha, \lambda^{-1}\mathbf{I})$ is a Gaussian likelihood with noise precision λ .² The relations between the different factors in the joint likelihood (4.21) are illustrated in Figure 4.4.

To obtain a computationally tractable estimator, we apply a variational EM algorithm to infer point estimates $\hat{\theta}$ and $\hat{\mathbf{z}}$ of θ and \mathbf{z} , respectively, while approximating the posterior distribution $p(\alpha, \gamma | \mathbf{y}, \theta, \mathbf{z})$ conditional on \mathbf{y} , θ and \mathbf{z} with a factorized distribution

$$q(\alpha, \gamma) = q_{\alpha}(\alpha) \prod_{k=1}^{K_{\max}} \prod_{l \in S_{\max}} q_{\gamma,k,l}(\gamma_{k,l}). \quad (4.22)$$

Estimates $\hat{\theta}$, $\hat{\mathbf{z}}$ and q are obtained by iteratively maximizing the ELBO

$$\mathcal{L}(q, \theta, \mathbf{z}) = \langle \ln p(\alpha, \theta, \gamma, \mathbf{z}, \mathbf{y}) - \ln q(\alpha, \gamma) \rangle_{q(\alpha, \gamma)} \quad (4.23)$$

with respect to q , θ and \mathbf{z} .

Update of α , θ , and \mathbf{z}

Instead of directly applying the variational M-step (2.52), $\hat{\theta}$ and $\hat{\mathbf{z}}$ are estimated jointly with the distribution $q_{\alpha}(\alpha)$ due to their strong interdependencies [E]. Maximizing the ELBO (4.23) jointly with respect to θ , \mathbf{z} and q_{α} results in the distribution

$$q_{\alpha}(\alpha; \hat{\theta}, \hat{\mathbf{z}}) = \text{CN}(\alpha_{\mathcal{S}}; \hat{\alpha}, \hat{\Sigma}) \prod_{\alpha_{k,l} \notin \alpha_{\mathcal{S}}} \delta(\alpha_{k,l}) \quad (4.24)$$

dependent on the estimates $\hat{\theta}$ and $\hat{\mathbf{z}}$, where $\alpha_{\mathcal{S}} = [\alpha_{k,l} : (k, l) \in \hat{\mathcal{S}}]^T$ is the vector comprised of amplitudes $\alpha_{k,l}$ with indices $(k, l) \in \hat{\mathcal{S}} \triangleq \{(k, l) : k \in \{1, 2, \dots, K_{\max}\}, l \in$

²For conciseness, λ is assumed known, whereas λ is assumed unknown and included in the estimation procedure in [E].

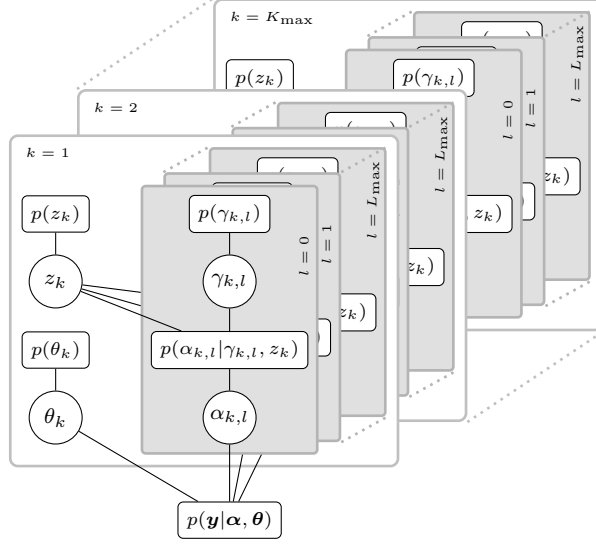


Figure 4.4: Factor graph of the Bernoulli-Gamma-Gaussian model for sparse group estimates assuming that $\mathcal{S}_{\max} = \{0, 1, \dots, L_{\max}\}$ for illustrative purposes. Sparsity of components is promoted by the hyperparameters $\gamma_{k,l}$, while sparsity in groups is promoted by the hyperparameters z_k .

$\mathcal{S}_{\max}, \hat{\gamma}_{k,l} < \infty, \hat{z}_k = 1\}$ corresponding to the “active” amplitudes estimated to be nonzero by the current estimates $\hat{\gamma}_{k,l} = \langle \gamma_{k,l} \rangle_{q_{\gamma,k,l}}$ and $\hat{\mathbf{z}} = [\hat{z}_1 \ \hat{z}_2 \ \dots \ \hat{z}_{K_{\max}}]^T$ of $\gamma_{k,l}$ and \mathbf{z} , respectively. The mean $\hat{\boldsymbol{\alpha}}$ and covariance $\hat{\boldsymbol{\Sigma}}$ of the normal distribution in $q_{\boldsymbol{\alpha}}$ are given by

$$\hat{\boldsymbol{\alpha}} = \lambda \hat{\boldsymbol{\Sigma}} \hat{\boldsymbol{\Psi}}^H \mathbf{y} \quad \hat{\boldsymbol{\Sigma}} = (\lambda \hat{\boldsymbol{\Psi}}^H \hat{\boldsymbol{\Psi}} + \hat{\boldsymbol{\Gamma}})^{-1} \quad (4.25)$$

where $\hat{\boldsymbol{\Psi}} = [e^{j2\pi g(\hat{\theta}_k, l)t} : (k, l) \in \hat{\mathcal{S}}]$ is a matrix with the columns of $\boldsymbol{\Psi}(\hat{\boldsymbol{\theta}})$ corresponding to the nonzero amplitudes $\boldsymbol{\alpha}_{\hat{\mathcal{S}}}$ parameterized by $\hat{\boldsymbol{\theta}}$, and $\hat{\boldsymbol{\Gamma}} = \text{diag}([\hat{\gamma}_{k,l} : (k, l) \in \hat{\mathcal{S}}]^T)$ is a matrix with the estimates of the prior precision $\hat{\gamma}_{k,l}$ of the “active” spectral lines of each source on its main diagonal.

Assuming an uninformative prior $p(\boldsymbol{\theta}) = \text{const.}$, the estimates $\hat{\boldsymbol{\theta}}$ and $\hat{\mathbf{z}}$ are obtained by coordinate ascent for one tuple of elements $(\hat{\theta}_k, \hat{z}_k)$ at a time in an iterative manner. Specifically, we show in [E] that if $\exists \theta_k : \Delta_k(\theta_k) > 0$, the ELBO (4.23) is maximized with respect to the tuple (θ_k, z_k) by the estimates $\hat{z}_k = 1$ and $\hat{\theta}_k = \arg \max_{\theta_k} \Delta_k(\theta_k)$, where the function $\Delta_k(\theta_k)$ is given by

$$\Delta_k(\theta_k) = \hat{\mathbf{u}}_k^H \hat{\boldsymbol{\Sigma}}_k^{-1} \hat{\mathbf{u}}_k + \ln |\hat{\boldsymbol{\Sigma}}_k| + \ln \frac{\bar{z}}{1 - \bar{z}} + \sum_{l \in \hat{\mathcal{S}}_k} (\ln \hat{\gamma}_{k,l} + \chi_0) \quad (4.26)$$

with

$$\hat{\mathbf{u}}_k = \lambda \hat{\boldsymbol{\Sigma}}_k \hat{\boldsymbol{\Psi}}_k^H (\mathbf{I} - \lambda \hat{\boldsymbol{\Psi}}_k \hat{\boldsymbol{\Sigma}}_k \hat{\boldsymbol{\Psi}}_k^H) \mathbf{y} \quad (4.27)$$

and

$$\hat{\boldsymbol{\Sigma}}_k = (\lambda \hat{\boldsymbol{\Psi}}_k^H \hat{\boldsymbol{\Psi}}_k + \hat{\boldsymbol{\Gamma}}_k - \lambda^2 \hat{\boldsymbol{\Psi}}_k^H \hat{\boldsymbol{\Psi}}_k \hat{\boldsymbol{\Sigma}}_k \hat{\boldsymbol{\Psi}}_k^H \hat{\boldsymbol{\Psi}}_k)^{-1} \quad (4.28)$$

implicitly depending on θ_k . If $\nexists \theta_k : \Delta_k(\theta_k) > 0$, the ELBO is maximized by $\hat{z}_k = 0$. In this case, the k th source is assumed to be nonexistent and, therefore, the ELBO (4.23) does not depend on θ_k . The variables $\hat{\mathbf{u}}_k$ and $\hat{\Sigma}_k$ in (4.27) and (4.28) correspond to the mean and covariance of the amplitudes of the k th source conditional on the current estimate of the remaining sources. In (4.27) and (4.28) $\hat{\Psi}_k \triangleq \hat{\Psi}_k(\theta_k) = [e^{j2\pi g(\theta_k, l)\mathbf{t}} : l \in \hat{\mathcal{S}}_k]$ are the dictionary entries corresponding to the spectral lines of the k th source parameterized with θ_k indexed by $\hat{\mathcal{S}}_k = \{l \in \mathcal{S}_{\max} : \hat{\gamma}_{k,l} < \infty\}$, and $\hat{\Psi}_{\bar{k}} = [e^{j2\pi g(\hat{\theta}_m, n)\mathbf{t}} : (m, n) \in \hat{\mathcal{S}}_{\bar{k}}]$ are the dictionary entries of the “active” spectral lines corresponding to all other sources, index by the set $\hat{\mathcal{S}}_{\bar{k}} = \{(m, n) \in \hat{\mathcal{S}} : m \neq k\}$. Furthermore, $\hat{\mathbf{I}}_k = \text{diag}([\hat{\gamma}_{k,l} : l \in \hat{\mathcal{S}}_k]^T)$, $\hat{\Sigma}_{\bar{k}} = (\lambda \hat{\Psi}_{\bar{k}}^H \hat{\Psi}_{\bar{k}} + \hat{\mathbf{I}}_{\bar{k}})^{-1}$, $\hat{\mathbf{I}}_{\bar{k}} = \text{diag}([\hat{\gamma}_{m,n} : (m, n) \in \hat{\mathcal{S}}_{\bar{k}}]^T)$ and, χ_0 is the digamma function evaluated at $\eta + 1$.

Update of γ

Solving the variational-EM update equation (2.53) for the distribution $q_{\gamma,k,l}$, we obtain $q_{\gamma,k,l}(\gamma_{k,l}) = p(\gamma_{k,l})$ if $\hat{z}_k = 0$, and

$$q_{\gamma,k,l}(\gamma_{k,l}) = \text{Ga}(\gamma_{k,l}; \eta + 1, \hat{\nu}_{k,l}) \quad (4.29)$$

otherwise, where $\hat{\nu}_{k,l} = \nu + \hat{\alpha}_{k,l} + \sigma_{k,l}^2$ and $\sigma_{k,l}^2$ is the element of the main diagonal of $\hat{\Sigma}$ that corresponds to the approximate (marginal) variance of the amplitude $\alpha_{k,l}$. The distributions $q_{\gamma,k,l}$ are fully described by their mean $\hat{\gamma}_{k,l} = \frac{\eta+1}{\hat{\nu}_{k,l}}$. We apply the fast update [D, 53], to find the mean estimates $\hat{\gamma}_{k,l}$ corresponding to repeated updates of q_{α} and $q_{\gamma,k,l}$ described in Section 4.2.2 as

$$\hat{\gamma}_{k,l} = \begin{cases} (|\hat{u}_{k,l}|^2 - \hat{s}_{k,l})^{-1} & \text{if } \frac{|\hat{u}_{k,l}|^2}{\hat{s}_{k,l}} > \chi_1 \\ \infty & \text{otherwise} \end{cases} \quad (4.30)$$

where the variables $\hat{s}_{k,l} = (\lambda \hat{\psi}_{k,l}^H \hat{\psi}_{k,l} - \lambda^2 \hat{\psi}_{k,l}^H \hat{\Psi}_{\bar{k},l} \hat{\Sigma}_{\bar{k},l} \hat{\Psi}_{\bar{k},l}^H \hat{\psi}_{k,l})^{-1}$ and $\hat{u}_{k,l} = \hat{s}_{k,l}(\lambda \hat{\psi}_{k,l}^H - \lambda^2 \hat{\psi}_{k,l}^H \hat{\Psi}_{\bar{k},l} \hat{\Sigma}_{\bar{k},l} \hat{\Psi}_{\bar{k},l}^H) \mathbf{y}$ are analogous to $\hat{\mathbf{u}}_k$ and $\hat{\Sigma}_k$ from above, except they are scalars since they characterize a single spectral line indexed by k and l , instead of jointly characterizing all lines within the k th group. In the definition of $\hat{u}_{k,l}$ and $\hat{s}_{k,l}$, the other variables are defined as $\hat{\psi}_{k,l} = e^{j2\pi g(\hat{\theta}_k, l)\mathbf{t}}$, $\hat{\Psi}_{\bar{k},l} = [e^{j2\pi g(\hat{\theta}_m, n)\mathbf{t}} : (m, n) \in \hat{\mathcal{S}} \setminus (k, l)]$, $\hat{\Sigma}_{\bar{k},l} = (\lambda \hat{\Psi}_{\bar{k},l}^H \hat{\Psi}_{\bar{k},l} + \hat{\mathbf{I}}_{\bar{k},l})^{-1}$, and $\hat{\mathbf{I}}_{\bar{k},l} = \text{diag}([\hat{\gamma}_{m,n} : (m, n) \in \hat{\mathcal{S}} \setminus (k, l)]^T)$. Finally, $\chi_1 \geq 1$ is an additional threshold that can be used to decrease the number of false alarms [E, 53, 60, 82].

Using the joint updates of $\hat{\theta}$, $\hat{\mathbf{z}}$ and q_{α} together with the updates of $q_{\gamma,k,l}$, a variational algorithm is presented in [E] for the estimation of the number of sources K , the parameters θ_k of each source, the index set $\mathcal{S}_k \forall k$ and the corresponding amplitudes $\alpha_{k,l}$. This algorithm starts with an empty model where $\hat{\mathbf{z}} = \mathbf{0}$ and all $\hat{\gamma}_{k,l} = \infty$, and iterates between searching for a new source to add to the model and refining (and possibly pruning) existing sources. See [E] for details about the algorithm and its implementation.

Thresholding for Groups and Individual Spectral Lines

The amplitude vector α is sparse on two levels: the number of groups and the number of spectral lines within each group. Naturally, the group-sparse super-

resolution algorithm [E] has two parameters that can be used to tune the sparsity achieved at each level. The threshold χ_1 in (4.30) can be used to tune the sparsity obtained within each group, i.e. to tune the detection threshold for individual spectral lines belonging to active sources, similar to the threshold for canonical (non-grouped) LSE derived in [60]. Furthermore, the parameter \bar{z} in the prior $p(\mathbf{z})$ can be used to tune the detection threshold for new sources.

The double sparsity introduces ambiguities. Instead of adding another spectral line to an already existing group, we can equivalently create a new group parameterized to result in a spectral line with the same frequency. E.g., if the individual frequencies for each group of spectral lines are given by $g(\theta_k, l) = \theta_k + \Delta l$ for some fixed spacing Δ , then instead of adding a spectral line with frequency $g(\theta_1, 1) = \theta_1 + \Delta$ to the group parameterized by θ_1 , we can introduce a new group parameterized with $\theta_2 \triangleq \theta_1 + \Delta$ consisting of a line indexed by $l = 0$, such that the frequency $g(\theta_2, 0) = \theta_2 = \theta_1 + \Delta$ is equivalent to the first case. However, if we add a new group then we also add a new parameter θ_2 which increases the degrees of freedom in the model. Using the principle of Occam's razor, we should seek the simplest description that fits the data, i.e. the model with the least degrees of freedom or, equivalently, the smallest number of groups K .

To promote models with a small number of groups, the addition of new groups to the model should be penalized more than the addition of new spectral lines to existing groups. To achieve this, we express the parameter \bar{z} of the independent priors $p(z_k) = \bar{z}^{z_k} (1 - \bar{z})^{(1-z_k)}$ as

$$\bar{z} = \frac{1}{1 + e^{(\chi_0 + \chi_2)}} \quad (4.31)$$

in terms of a second threshold χ_2 . In [E] we show that by choosing $\chi_2 > \chi_1$, the addition of new groups with a single spectral line is penalized more than the addition of new spectral lines to already existing groups.

4.3.2 Application: Extended Object Detection

We perform a numerical study to assess the performance of the algorithm proposed in [E] for the detection of extended objects. To highlight the gain of considering the structural relations between the spectral lines, we simulate a single extended object in low-SNR conditions. The radar return from the object is modeled as a multipath channel $h_k(t)$ as (2.11). The delays $\tau_{k,l}$ of these paths are drawn from a point process with uniform intensity function between $32.3232T$ and $37.323T$ such that on average 10 MPCs from the object are considered, where $T = \frac{1}{B}$ is the Nyquist-sampling frequency of the system with bandwidth B . The corresponding amplitudes $\alpha_{k,l}$ are drawn from a zero-mean proper complex Gaussian distribution with unit variance.

Since no other algorithm directly integrates the estimation of the extent of the object into the LSE problem, we compare the proposed algorithm in a setup intended to reflect a tracking scenario in which the LSE algorithm is used as pre-processing step for a multi-object tracking algorithm, e.g. based on belief propagation and joint probabilistic data association [86]. Specifically, we use an oracle data association, which associates all components within $\pm T$ of the true object region to belong to the object, reflecting the association capabilities of a tracking filter applied to the




Figure 4.5: Comparison of the algorithm proposed in [E] compared to the FV-SBL [47], SF-SBL [85] and NOMP [22]. Subfigure (a) shows the cumulative frequency of the RMSE at an SNR of -6 dB, (b) shows the mean RMSE over SNR, and (c) shows the histogram of the delay of the detected components at an SNR of -6 dB.

measurements in post-processing. From the oracle data association, we compute the estimated extent

$$\hat{O}_E = \max_l \hat{\tau}_{o,l} - \min_l \hat{\tau}_{o,l} \quad (4.32)$$

and center of mass

$$\hat{O}_C = \frac{\sum_{l=1}^{\hat{L}} \hat{\alpha}_{o,l}^2 \hat{\tau}_{o,l}}{\sum_{l=1}^{\hat{L}} \hat{\alpha}_{o,l}^2} \quad (4.33)$$

of the object, where \hat{L} is the number of spectral lines associated with the object by the oracle data association, and $\hat{\tau}_{o,l}$ and $\hat{\alpha}_{o,l}$, $l \in \{1, 2, \dots, \hat{L}\}$ are the delays and amplitudes of the associated spectral lines, respectively. As performance metric we use the root mean square error (RMSE), computed as

$$\text{RMSE} = \sqrt{(\hat{O}_E - O_E)^2 + (\hat{O}_C - O_C)^2} \quad (4.34)$$

where O_E and O_C are the true extent and center of mass of the object.

As comparison algorithms we use the Newtonized orthogonal matching pursuit (NOMP) LSE algorithm [22], as well as two variants of (non-grouped) SBL-based

LSE algorithms, the superfast SBL (SF-SBL) [85] and fast variational SBL (FV-SBL) [47]. For a fair comparison, all algorithms are parameterized such that they produce approximately the same rate of erroneous detection outside the object region. Figure 4.5a shows the cumulative frequency of the RMSE at an SNR of $\frac{\|y-w\|^2}{\|w\|^2} = -6$ dB whereas Figure 4.5b shows the mean RMSE as a function of the SNR. All three comparison methods achieve virtually the same performance. However, the proposed algorithm achieves a smaller mean RMSE for all simulated SNRs and a higher cumulative frequency for smaller RMSEs at SNR = -6 dB. Figure 4.5c shows a histogram of the delays estimated by the different algorithms, which illustrates the reason for this performance gain. Outside the object region, all algorithms exhibit a similar false alarm rate. However, inside the object region, the proposed algorithm is able to detect more spectral lines than the comparison methods. This is due to the underlying structured LSE model (4.1) which assumes spectral lines appear in closely-spaced groups. Once a spectral line is detected in the object region, the grouped model increases the probability of detection for other spectral lines in the vicinity, resulting in a performance gain when estimating the parameters of the extended object.

4.3.3 Application: Multi-Pitch Estimation

The BSURE-IR algorithm for multipitch estimation of audio data [84] is the only other algorithm found in the literature which solves the same class of problems as [E], i.e. estimating the number of sources K and the index sets \mathcal{S}_k of each source in a structured line spectrum (4.1), together with estimating the group parameters θ_k on a continuum. To compare both algorithms, we apply the algorithm presented in [E] to the multi-pitch estimation problem.

Consider K tones being produced by multiple instruments to form a chord within a musical piece. The aim of multi-pitch estimation is to estimate the fundamental frequencies f_k and the number of sources K based on the audio signal $y(t)$ being recorded by a microphone. Assuming each source produces a stationary periodic signal with fundamental frequency (i.e. pitch) f_k , the signal $y(t)$ is modeled as

$$y(t) = \sum_{k=1}^K \sum_{l \in \mathcal{S}_k} \alpha_{k,l} \cdot e^{j2\pi l f_k t} + w(t) \quad (4.35)$$

where K is the number of tones played simultaneously, $\mathcal{S}_k \subset \mathbb{N}$ is a set of harmonics produced by the instrument playing the k th tone, $\alpha_{k,l}$ is the amplitude of the l th harmonic of the k th tone in the chord, and $w(t)$ is the measurement noise [87]. See Box 4.2 for how to obtain the complex-valued signal model (4.35) from real-valued audio data. Equation (4.35) is an instance of the structured LSE model (4.1), where $g(\theta_k, l) = l\theta_k$.

The performance of the algorithm proposed in [E] is compared against the BSURE-IR algorithm using numerical simulations (see [E] for details) and on real audio data using the Bach 10 dataset [88]. The dataset contains 10 musical pieces composed by J. S. Bach played by a quartet consisting of a violin, clarinet, saxophone, and bassoon. Each instrument was recorded individually while listening to the others via headphones. The individual recordings are mixed together to form the recordings contained in the dataset. The ground truth is extracted from

Box 4.2: Complex-Valued Signal Model for Real-Valued Audio Data.

Audio data of a microphone is typically real-valued, since it corresponds to the pressure of the sound waves at a point in space. For the application of multi-pitch estimation, this data can be modeled by

$$y(t) = \sum_{k=1}^K \sum_{l \in \mathcal{S}_k} \alpha_{k,l} \cos(2\pi l f_k t + \varphi_{k,l}) + w(t) \quad (4.36)$$

where $\varphi_{k,l}$ is the phase and $\alpha_{k,l} \in \mathbb{R}_{\geq 0}$ the real-valued amplitude of the l th harmonic of the k th source. Since it is typically not possible to marginalize over the phases $\varphi_{k,l}$ due to the nonlinear dependency of $y(t)$ on $\varphi_{k,l}$, they must be included as nuisance parameters in any LSE algorithm based on such a real-valued signal model.

To sidestep the problem of estimating the phases $\varphi_{k,l}$, the signal $y(t)$ can be transformed to a complex-valued equivalent by computing the analytic signal $y_A(t) = y(t) + j\hat{y}(t)$, where $\hat{y}(t)$ is the Hilbert-transform of $y(t)$. This results in a complex valued signal $y_A(t)$ containing only the positive frequencies of the Hermitian symmetric spectrum of $y(t)$, which can be modeled using (4.35).

the individual recordings using the YIN single pitch detection algorithm [89]. For processing, the recordings are segmented into frames of length 45 ms with a 10 ms stride between frames. Each frame was processed individually and each detected fundamental frequency considered “matched” if it deviated no more than half a semitone (approximately 3%) from a pitch in the ground truth. We compare the algorithms using metrics of accuracy, precision, recall and F_1 score defined as

$$\text{Accuracy} = \frac{\text{TP}}{\text{TP} + \text{FP} + \text{FN}} \quad (4.37)$$

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad (4.38)$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (4.39)$$

$$F_1 = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (4.40)$$

where TP is be the number of true positives, i.e. the fundamental frequencies matched between the ground truth and the estimate in all frames, FP is the number of false positives, i.e. estimated fundamental frequencies that did not match to any fundamental frequencies in the ground truth, and FN is the number of false negatives, i.e. ground truth fundamental frequencies that did not match any estimated frequencies. Table 4.1 lists the accuracy, precision, recall and the F_1 score of the proposed algorithm along with several comparison methods. In addition to the algorithm proposed in [E] and the BSURE-IR algorithm [84], Table 4.1 also includes a data-driven method pretrained on the instruments in the dataset denoted as “BW15” [90], and another model-driven method denoted PEARLS [91], which estimates the fundamental frequencies on a grid. The best performance in both precision and recall is achieved by the method proposed in [E]. Since accuracy and F_1 score are compound measures depending on precision and recall, the proposed

Table 4.1: Multi-Pitch estimation performance of [E] evaluated on the Bach-10 dataset.

Method	F_1	Accuracy	Precision	Recall	Pre-Trained
Prop. Alg.	0.72	0.56	0.73	0.70	No
BW15	0.67	0.52	0.68	0.68	Yes
BSURE-IR	0.64	0.47	0.68	0.54	No
PEARLS	0.60	0.44	0.56	0.51	No

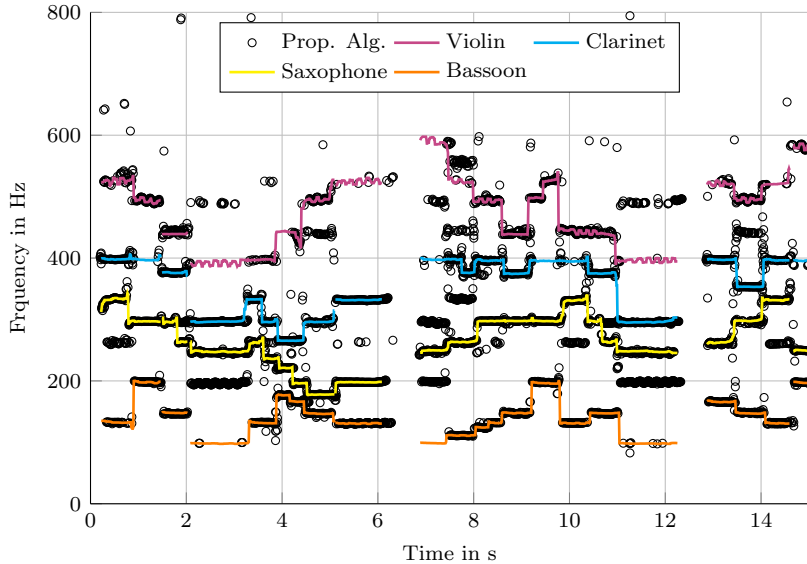


Figure 4.6: Fundamental frequency estimates for 15 seconds of the choral *Ach Gott und Herr* compared to the ground truth (best viewed in color). The performance of the algorithm proposed in [E] for the depicted timespan is accuracy = 0.55, precision = 0.75 and recall = 0.68.

algorithm performs best according to those measures as well.

Fig 4.6 illustrates the pitches detected by the algorithm proposed in [E] for the first 15 seconds of the piece “Ach Gott und Herr” from the dataset. Some errors can be seen, e.g. between approximately 2s and 3.5s, the bassoon is not detected at all and the fundamental frequency of the violin is estimated to be half the true frequency (octave error). Some spurious artifacts can also be seen throughout the plot. Nevertheless, overall the proposed algorithm is able to estimate the pitches quite accurately. E.g., the saxophone is detected with near-perfect accuracy. Furthermore, the evaluation over the whole dataset presented in Table 4.1 demonstrates that the algorithm proposed in [E] is able to outperform methods developed specifically for multi-pitch estimation, such as BSURE-IR [84] and PEARLS [91], as well as data-driven methods that have been pre-trained on the instruments in the dataset [90].

4.4 Concluding Remarks

This chapter investigates Hypothesis 2, i.e. the hypothesis that the performance of sparsity-based LSE algorithms such as SBL can be improved for the estimation of structured line spectra using group-sparsity. We introduce the structured LSE model (4.1), which arises naturally in multiple applications, e.g. the detection of extended objects. Two SBL-based LSE algorithms are presented which utilize group-sparsity to infer the parameters of such structured line spectra. Assuming known group sizes, the “fast variational block SBL” algorithm presented in [D] is shown to significantly improve the runtime compared to other SBL-based algorithms using group-sparsity, such as BSBL [28], while also achieving a slightly better estimation performance in terms of the NMSE of the amplitudes. Assuming unknown group sizes, the “group-sparse super-resolution LSE” algorithm presented in [E] is shown to outperform other LSE methods for the detection and estimation of extended objects. Analysis indicates that the performance gain is due to the clustering of the spectral lines into related groups, i.e. by considering the structural relation of the lines in the spectrum. Doing so lowers the detection threshold in the vicinity of spectral lines which have already been detected, increasing the detection probability of other spectral lines likely belonging to the same extended object. Thus, the number of spectral lines detected for each object in low-SNR conditions is increased. As a result, the the objects parameters, such as its extent and center of mass, can be estimated more accurately.

LSE is a fundamental problem in many domains of science and engineering. Naturally, the grouped LSE model (4.1) applies to other domains as well, allowing the algorithms developed in [D, E] to be applied to different types of data as well. Sections 4.2.4 and 4.3.3 illustrates two such examples. Specifically, we apply the fast variational block SBL algorithm [D] to DOA estimation based on MMVs received from a microphone or antenna array, and we use the group-sparse super-resolution algorithm [E] for multi-pitch detection of audio signals. In both cases, we demonstrate superior performance compared to the state-of-the-art methods, demonstrating the versatility of the developed algorithms.

5

Conclusion

Two research questions concerning the detection and estimation of dispersive target signals are investigated in this thesis. The answers to the research hypotheses defined at the beginning of this thesis are summarized as follows, based on the results presented in sections 3 and 4.

Hypothesis 1: *The performance of car occupancy detection based on UWB radar can be improved by modeling and estimating the dispersive multipath channel between the transmitter, the target, and the receiver.*

In short, Hypothesis 1 is shown to be true. Both of the model-based methods presented in [A,B] as well as the data-driven method presented in [C] outperform the baseline approaches of the energy detector (ED) or FFT detector, which do not account for the multipath propagation between the transmitter, target, and receiver. The model-based methods are obtained by modeling the human body as a point-like target in order to apply a pinhole channel model. Although this assumption may be somewhat unrealistic, the detection algorithms based on this model outperform other model-based approaches for car occupancy detection using UWB radar. Hence, in the spirit of George E. P. Box [1], we consider the model obtained by this simplification to be useful. The data-driven method [C] outperforms the model-based methods [A,B] presented in this thesis. There are several possible reasons why the data-driven method outperforms the model-based methods. E.g., due to model mismatch, or because the true distribution of the random variables differs from the comparatively simple prior distributions used in [A,B]. The performance of the presented car occupancy detection methods is evaluated on the UWBCarGraz dataset [68], which consists of data from adult participants only. Before deploying any of the car occupancy detection algorithms presented in [A–C] into cars, the performance of these algorithms should be verified on cases not included in the UWBCarGraz dataset, e.g. infants, small children, or multiple people occupying the car at the same time. Due to their model-based approach, the algorithms presented in [A,B] can be expected to generalize better to such cases than the purely data-driven approach presented in [C].

Hypothesis 2: *Group sparsity can be used to generalize sparsity-based LSE methods to the case of structured line spectra. Thereby, the performance of LSE algorithms can be improved for applications where such structured line spectra arise naturally.*

The results presented in [D,E], which are summarized in Chapter 4, confirm this hypothesis. E.g., the analysis presented in [E] shows that the group-sparse super-resolution LSE algorithm outperforms LSE algorithms that are based on the canonical assumption of uncorrelated spectral lines when detecting a single extended object in low-SNR conditions. The algorithms developed in [D,E] are also shown to outperform state-of-the-art algorithms for multi-pitch estimation [84], and DOA estimation based on MMVs [54] in their, respective, application domains. The superior performance in the case of multi-pitch estimation is particularly noteworthy because the BSURE algorithm [84] is the only algorithm found in the literature for the estimation of structured line spectra considering unknown group-sizes which estimates the group parameters on a continuum. I.e., it is the only algorithm solving the exact same inference problem as the group-sparse super-resolution algorithm presented in [E].

An interesting perspective for future research is to identify other application domains where structured line spectra naturally arise, in order to investigate any potential performance gain from applying the algorithms presented in [D,E]. E.g., (variational) SBL has been successfully integrated into the joint channel estimation and data detection problem by merging belief propagation on factor graphs with variational Bayesian inference [92,93]. Channels of multiple-input multiple-output systems are often found to be approximately block-sparse, thus, group-sparse variants of SBL are used for joint channel estimation and data detection in these problems [94]. This is a promising application area where the computational efficiency could be improved by using the fast variational block SBL algorithm [D], or the accuracy of the obtained channel estimate could be improved by using the group-sparse super-resolution algorithm [E]. Finally, the columns of the dictionary matrix can also be constructed from (parameterized) functions other than complex exponentials, i.e. spectral lines, as well. Thus, the developed algorithms [D,E] can, in theory, be applied to an even wider range of applications where the signal is modeled as a superposition of groups of correlated components.

Bibliography

- [1] G. E. P. Box, “Robustness in the strategy of scientific model building,” in *Robustness in Statist.*, Apr. 11–12, 1978, doi: 10.1016/B978-0-12-438150-6.50018-2.
- [2] H.-J. Pirch and F. Leong. (2020) Introduction to impulse radio UWB seamless access systems. FiRa Consortium. Accessed: Mar. 2024. [Online]. Available: <https://www.firaconsortium.org/resource-hub/white-papers>
- [3] A. F. Molisch, “Ultra-wide-band propagation channels,” *Proc. IEEE*, Feb. 2009, doi: 10.1109/JPROC.2008.2008836.
- [4] *IEEE Standard for Low-Rate Wireless Networks*, IEEE Std. 802.15.4-2020, Jul. 2020, (Revision of IEEE Std 802.15.4-2015).
- [5] S. Ahmed and S. H. Cho, “Machine learning for healthcare radars: Recent progresses in human vital sign measurement and activity recognition,” *IEEE Commun. Surveys Tuts.*, 2023, doi: 10.1109/COMST.2023.3334269.
- [6] Y. Kilic, H. Wymeersch, A. Meijerink, M. J. Benthum, and W. G. Scanlon, “Device-free person detection and ranging in UWB networks,” *IEEE J. Sel. Topics Signal Process.*, Feb. 2014, doi: 10.1109/JSTSP.2013.2281780.
- [7] S. Yang, J. L. Kerneq, O. Romain, F. Fioranelli, P. Cadart, J. Fix, C. Ren, G. Manfredi, T. Letertre, I. D. H. Sáenz, J. Zhang, H. Liang, X. Wang, G. Li, Z. Chen, K. Liu, X. Chen, J. Li, X. Wu, Y. Chen, and T. Jin, “The human activity radar challenge: Benchmarking based on the ‘radar signatures of human activities’ dataset from glasgow university,” *IEEE J. Biomed. Health Inform.*, Apr. 2023, doi: 10.1109/JBHI.2023.3240895.
- [8] D. Yang, Z. Zhu, J. Zhang, and B. Liang, “The overview of human localization and vital sign signal measurement using handheld IR-UWB through-wall radar,” *Sensors*, May 2021, doi: 10.3390/s21020402.
- [9] A. S. Aghaei, B. Donmez, C. C. Liu, D. He, G. Liu, K. N. Plataniotis, H.-Y. W. Chen, and Z. Sojoudi, “Smart driver monitoring: When signal processing meets human factors: In the driver’s seat,” *IEEE Signal Process. Mag.*, Nov. 2016, doi: 10.1109/MSP.2016.2602379.
- [10] B. Schleicher, I. Nasr, A. Trasser, and H. Schumacher, “IR-UWB radar demonstrator for ultra-fine movement detection and vital-sign monitoring,” *IEEE Trans. Microw. Theory Techn.*, May 2013, doi: 10.1109/TMTT.2013.2252185.

- [11] S. Lim, S. Lee, J. Jung, and S.-C. Kim, "Detection and localization of people inside vehicle using impulse radio ultra-wideband radar sensor," *IEEE Sensors J.*, Apr. 2020, doi: 10.1109/JSEN.2019.2961107.
- [12] Y. Ma, Y. Zeng, and V. Jain, "CarOSense: Car occupancy sensing with the ultra-wideband keyless infrastructure," *Proc. ACM Interact. Mobile Wearable Ubiquitous Technol.*, Sep. 2020, doi: 10.1145/3411820.
- [13] Euro NCAP. (2021) Euro NCAP 2025 roadmap. Accessed: May 2021. [Online]. Available: <https://cdn.euroncap.com/media/30701/euroncap-roadmap-2025-v4-print.pdf>
- [14] H. L. Van Trees, *Detection, Estimation, and Modulation Theory Part III: Radar-Sonar Processing and Gaussian Signals in Noise*. Hoboken, NJ, USA: Wiley, 2001, doi: 10.1002/0471221090.
- [15] I. M. Skolnik, *Radar Handbook*, 3rd ed. New York, NY, USA: McGraw-Hill, 2008.
- [16] K. Granstrom, M. Baum, and S. Reuter, "Extended object tracking: Introduction, overview, and applications," *J. Adv. Inf. Fusion*, vol. 12, no. 2, pp. 139–174, Dec. 2017.
- [17] F. Bellens, F. Quitin, F. Horlin, and P. De Doncker, "UWB channel analysis within a moving car," in *9th Int. Conf. Intell. Transp. Syst. Telecommun.*, Oct. 20–22, 2009, doi: 10.1109/ITST.2009.5399271.
- [18] O. Delangre, S. Van Roy, P. De Doncker, M. Lienard, and P. Degauque, "Modeling in-vehicle wideband wireless channels using reverberation chamber theory," in *IEEE 66th Veh. Technol. Conf.*, Sep. 30–Oct. 3, 2007, doi: 10.1109/VETECF.2007.451.
- [19] Y. Katayama, K. Terasaka, K. Higashikaturagi, I. Matunami, and A. Kajiwara, "Ultra-wideband impulse-radio propagation for in-vehicle wireless link," in *IEEE 64th Veh. Technol. Conf.*, Sep. 25–28, 2006, doi: 10.1109/VTCF.2006.63.
- [20] T. Kobayashi, "Measurements and characterization of ultra wideband propagation channels in a passenger-car compartment," in *IEEE 9th Int. Symp. Spread Spectr. Techn. and Appl.*, Aug. 28–31, 2006, doi: 10.1109/ISSSTA.2006.311768.
- [21] T. L. Hansen, "Sparsity-based algorithms for line spectral estimation," PhD Thesis, Aalborg University, Mar. 2018, doi: 10.5278/vbn.phd.tech.00037.
- [22] B. Mamandipoor, D. Ramasamy, and U. Madhow, "Newtonized orthogonal matching pursuit: Frequency estimation over the continuum," *IEEE Trans. Signal Process.*, Oct. 2016, doi: 10.1109/TSP.2016.2580523.
- [23] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Rev.*, Mar. 2001, doi: 10.1137/S003614450037906X.
- [24] K. Kreutz-Delgado, J. F. Murray, B. D. Rao, K. Engan, T.-W. Lee, and T. J. Sejnowski, "Dictionary learning algorithms for sparse representation," *Neural Comput.*, Feb. 2003, doi: 10.1162/089976603762552951.

- [25] E. J. Candés, M. B. Wakin, and S. P. Boyd, “Enhancing sparsity by reweighted ℓ_1 minimization,” *J. Fourier Anal. and Appl.*, Dec. 2008, doi: 10.1007/s00041-008-9045-x.
- [26] I. Gorodnitsky and B. Rao, “Sparse signal reconstruction from limited data using FOCUSS: a re-weighted minimum norm algorithm,” *IEEE Trans. Signal Process.*, Mar. 1997, doi: 10.1109/78.558475.
- [27] Z. Zhang and B. D. Rao, “Sparse signal recovery with temporally correlated source vectors using sparse Bayesian learning,” *IEEE J. Sel. Topics Signal Process.*, 2011, doi: 10.1109/JSTSP.2011.2159773.
- [28] —, “Extension of SBL algorithms for the recovery of block sparse signals with intra-block correlation,” *IEEE Trans. Signal Process.*, Apr. 2013, doi: 10.1109/TSP.2013.2241055.
- [29] C. M. Bishop, *Pattern Recognition and Machine Learning*, ser. Information Science and Statistics. Secaucus, NJ, USA: Springer-Verlag, New York, 2006.
- [30] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*, A. V. Oppenheim, Ed. Englewood Cliffs, NJ, USA: Prentice-Hall, 1993.
- [31] A. F. Molisch, *Wireless Communications*, 2nd ed. Hoboken, NJ, USA: Wiley, 2011.
- [32] N. Michelusi, U. Mitra, A. F. Molisch, and M. Zorzi, “UWB sparse/diffuse channels, part I: Channel models and Bayesian estimators,” *IEEE Trans. Signal Process.*, Jun. 2012, doi: 10.1109/TSP.2012.2205681.
- [33] P. Stoica and R. Moses, *Spectral analysis of signals*. Upper Saddle River, NJ, USA: Pearson Prentice Hall, 2005.
- [34] R. Schmidt, “Multiple emitter location and signal parameter estimation,” *IEEE Trans. Antennas Propag.*, Mar. 1986, doi: 10.1109/TAP.1986.1143830.
- [35] R. Roy and T. Kailath, “ESPRIT-estimation of signal parameters via rotational invariance techniques,” *IEEE Trans. Acoust., Speech, Signal Process.*, Jul. 1989, doi: 10.1109/29.32276.
- [36] I. Ziskind and M. Wax, “Maximum likelihood localization of multiple sources by alternating projection,” *IEEE Trans. Acoust., Speech, Signal Process.*, Oct. 1988, doi: 10.1109/29.7543.
- [37] M. Feder and E. Weinstein, “Parameter estimation of superimposed signals using the EM algorithm,” *IEEE Trans. Acoust., Speech, Signal Process.*, Apr. 1988, doi: 10.1109/29.1552.
- [38] H. Akaike, “A new look at the statistical model identification,” *IEEE Trans. Autom. Control*, Dec. 1974, doi: 10.1109/TAC.1974.1100705.
- [39] G. Schwarz, “Estimating the dimension of a model,” *Ann. Statist.*, vol. 6, no. 2, pp. 461–464, Mar. 1978.

- [40] P. Stoica and Y. Selen, “Model-order selection: a review of information criterion rules,” *IEEE Signal Process. Mag.*, Jul. 2004, doi: 10.1109/MSP.2004.1311138.
- [41] R. Tibshirani, “Regression shrinkage and selection via the LASSO,” *J. Roy. Statistical Soc.: Ser. B (Statistical Methodology)*, 1996, doi: 10.1111/j.2517-6161.1996.tb02080.x.
- [42] S. Mallat and Z. Zhang, “Matching pursuits with time-frequency dictionaries,” *IEEE Trans. Signal Process.*, 1993, doi: 10.1109/78.258082.
- [43] M. E. Tipping, “Sparse Bayesian learning and the relevance vector machine,” *J. Mach. Learn. Res.*, vol. 1, pp. 211–244, Jun. 2001.
- [44] M. E. Tipping and A. C. Faul, “Fast marginal likelihood maximisation for sparse Bayesian models,” in *Proc. 9th Int. Workshop Artif. Intell. and Statist.*, vol. R4, Key West, FL, USA, Jan. 03–06, 2003, pp. 276–283.
- [45] D. P. Wipf and B. D. Rao, “Sparse Bayesian learning for basis selection,” *IEEE Trans. Signal Process.*, Aug. 2004, doi: 10.1109/TSP.2004.831016.
- [46] T. L. Hansen, M. A. Badiu, B. H. Fleury, and B. D. Rao, “A sparse Bayesian learning algorithm with dictionary parameter estimation,” in *2014 IEEE 8th Sensor Array and Multichannel Signal Process. Workshop (SAM)*, Jun. 22–25, 2014, doi: 10.1109/SAM.2014.6882422.
- [47] D. Shutin, W. Wand, and T. Jost, “Incremental sparse Bayesian learning for parameter estimation of superimposed signals,” in *10th Int. Conf. Sampling Theory and Appl.*, Bremen, Germany, Jul. 1–5, 2013, pp. 513–516.
- [48] M.-A. Badiu, T. L. Hansen, and B. H. Fleury, “Variational Bayesian inference of line spectra,” *IEEE Trans. Signal Process.*, May 2017, doi: 10.1109/TSP.2017.2655489.
- [49] P. Stoica, P. Babu, and J. Li, “SPICE: A sparse covariance-based estimation method for array processing,” *IEEE Trans. Signal Process.*, Feb. 2011, doi: 10.1109/TSP.2010.2090525.
- [50] Z. Yang and L. Xie, “On gridless sparse methods for line spectral estimation from complete and incomplete data,” *IEEE Trans. Signal Process.*, Jun. 2015, doi: 10.1109/TSP.2015.2420541.
- [51] B. N. Bhaskar, G. Tang, and B. Recht, “Atomic norm denoising with applications to line spectral estimation,” *IEEE Trans. Signal Process.*, Dec. 2013, doi: 10.1109/TSP.2013.2273443.
- [52] Z. Yang and L. Xie, “Enhancing sparsity and resolution via reweighted atomic norm minimization,” *IEEE Trans. Signal Process.*, Feb. 2016, doi: 10.1109/TSP.2015.2493987.
- [53] D. Shutin, T. Buchgraber, S. R. Kulkarni, and H. V. Poor, “Fast variational sparse Bayesian learning with automatic relevance determination for superimposed signals,” *IEEE Trans. Signal Process.*, Dec. 2011, doi: 10.1109/TSP.2011.2168217.

- [54] P. Gerstoft, C. F. Mecklenbräuker, A. Xenaki, and S. Nannuru, “Multisnapshot sparse Bayesian learning for DOA,” *IEEE Signal Process. Lett.*, Oct. 2016, doi: 10.1109/LSP.2016.2598550.
- [55] M. F. Duarte and R. G. Baraniuk, “Spectral compressive sensing,” *Appl. Comput. Harmon. Anal.*, Jul. 2013, doi: 10.1016/j.acha.2012.08.003.
- [56] Y. Chi, L. L. Scharf, A. Pezeshki, and A. R. Calderbank, “Sensitivity to basis mismatch in compressed sensing,” *IEEE Trans. Signal Process.*, May 2011, doi: 10.1109/TSP.2011.2112650.
- [57] D. P. Wipf, B. D. Rao, and S. Nagarajan, “Latent variable Bayesian models for promoting sparsity,” *IEEE Trans. Inf. Theory*, Sep. 2011, doi: 10.1109/TIT.2011.2162174.
- [58] R. Chartrand and W. Yin, “Iteratively reweighted algorithms for compressive sensing,” in *2008 IEEE Int. Conf. Acoust., Speech and Signal Process.*, Mar. 31–Apr. 04, 2008, doi: 10.1109/ICASSP.2008.4518498.
- [59] J. O. Berger, *Statistical Decision Theory and Bayesian Analysis*, 2nd ed. New York, NY, USA: Springer-Verlag, 1980.
- [60] S. Grebien, E. Leitinger, K. Witrisal, and B. H. Fleury, “Super-resolution estimation of UWB channels including the dense component – an SBL-inspired approach,” *IEEE Trans. Wireless Commun.*, Feb. 2024, doi: 10.1109/TWC.2024.3371352.
- [61] D. J. C. MacKay, “Bayesian interpolation,” *Neural Comput.*, May 1992, doi: 10.1162/neco.1992.4.3.415.
- [62] M. Tipping, “The relevance vector machine,” in *Adv. Neural Inf. Process. Syst.*, vol. 12. Denver, CO, USA: MIT Press, Nov. 29 – Dec. 4, 1999, pp. 652–658.
- [63] J. Palmer, B. Rao, and D. Wipf, “Perspectives on sparse Bayesian learning,” in *Adv. Neural Inf. Process. Syst.*, vol. 16, Vancouver, Canada, Dec. 8–13, 2003, pp. 249–256.
- [64] J. Palmer, K. Kreutz-Delgado, B. Rao, and D. Wipf, “Variational EM algorithms for non-Gaussian latent variable models,” in *Adv. Neural Inf. Process. Syst.*, vol. 18, Vancouver, Canada, Dec. 5–8, 2005, pp. 1059–1066.
- [65] J. M. Winn, “Variational message passing and its applications,” PhD Thesis, University of Cambridge, Jan. 2004.
- [66] D. G. Tzikas, A. C. Likas, and N. P. Galatsanos, “The variational approximation for Bayesian inference,” *IEEE Signal Process. Mag.*, Nov. 2008, doi: 10.1109/MSP.2008.929620.
- [67] J. Winn, C. M. Bishop, and T. Jaakkola, “Variational message passing,” *J. Mach. Learn. Res.*, vol. 6, no. 4, pp. 661–694, Apr. 2005.
- [68] J. Möderl, S. Posch, F. Pernkopf, and K. Witrisal, “UWBCarGraz dataset,” Graz University of Technology, Nov. 2023, doi: 10.3217/2gx2m-pt043.

- [69] D. Arnitz, U. Muehlmann, and K. Witrisal, "Wideband characterization of backscatter channels: Derivations and theoretical background," *IEEE Trans. Antennas Propag.*, Jan. 2012, doi: 10.1109/TAP.2011.2167923.
- [70] S. M. Kay, *Fundamentals of Statistical Signal Processing: Detection Theory*, 14th ed., A. V. Oppenheim, Ed. Upper Saddle River, NJ, USA: Prentice-Hall, 1998.
- [71] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *2016 IEEE Conf. on Comput. Vision and Pattern Recognit.*, Jun. 27–30, 2016, doi: 10.1109/CVPR.2016.90.
- [72] M. Yuan and Y. Lin, "Model selection and estimation in regression with grouped variables," *J. Roy. Statistical Soc.: Ser. B (Statistical Methodology)*, Feb. 2006, doi: 10.1111/j.1467-9868.2005.00532.x.
- [73] M. Kyung, J. Gill, M. Ghosh, and G. Casella, "Penalized regression, standard errors, and Bayesian LASSOs," *Bayesian Anal.*, Jun. 2010, doi: 10.1214/10-BA607.
- [74] X. Xu and M. Ghosh, "Bayesian variable selection and estimation for group LASSO," *Bayesian Anal.*, Dec. 2015, doi: 10.1214/14-BA929.
- [75] Y. C. Eldar, P. Kuppinger, and H. Bolcskei, "Block-sparse signals: Uncertainty relations and efficient recovery," *IEEE Trans. Signal Process.*, Jun. 2010, doi: 10.1109/TSP.2010.2044837.
- [76] T. Kronvall, S. I. Adalbjörnsson, S. Nadig, and A. Jakobsson, "Group-sparse regression using the covariance fitting criterion," *Signal Process.*, Oct. 2017, doi: 10.1016/j.sigpro.2017.03.025.
- [77] D. P. Wipf and B. D. Rao, "An empirical Bayesian strategy for solving the simultaneous sparse approximation problem," *IEEE Trans. Signal Process.*, Jun. 2007, doi: 10.1109/TSP.2007.894265.
- [78] M. Luessi, S. D. Babacan, R. Molina, and A. K. Katsaggelos, "Bayesian simultaneous sparse approximation with smooth signals," *IEEE Trans. Signal Process.*, Nov. 2013, doi: 10.1109/TSP.2013.2280441.
- [79] Z. Ma, W. Dai, Y. Liu, and X. Wang, "Group sparse Bayesian learning via exact and fast marginal likelihood maximization," *IEEE Trans. Signal Process.*, May 2017, doi: 10.1109/TSP.2017.2675867.
- [80] S. D. Babacan, S. Nakajima, and M. N. Do, "Bayesian group-sparse modeling and variational inference," *IEEE Trans. Signal Process.*, Jun. 2014, doi: 10.1109/TSP.2014.2319775.
- [81] B. Jorgensen, *Statistical properties of the generalized inverse Gaussian distribution*, ser. Lecture notes in statistics. New York, NY, USA: Springer-Verlag, 1982, doi: 10.1007/978-1-4612-5698-4.

- [82] E. Leitinger, S. Grebien, B. Fleury, and K. Witrisal, "Detection and estimation of a spectral line in MIMO systems," in *2020 54th Asilomar Conf. Signals, Syst. and Computers*, Nov. 01–04, 2020, doi: 10.1109/IEEECONF51394.2020.9443381.
- [83] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, "A consistent metric for performance evaluation of multi-object filters," *IEEE Trans. Signal Process.*, Aug. 2008, doi: 10.1109/TSP.2008.920469.
- [84] J. Swärd, H. Li, and A. Jakobsson, "Off-grid fundamental frequency estimation," *IEEE/ACM Trans. Audio, Speech, Language Process.*, Feb. 2018, doi: 10.1109/TASLP.2017.2775800.
- [85] T. L. Hansen, B. H. Fleury, and B. D. Rao, "Superfast line spectral estimation," *IEEE Trans. Signal Process.*, Feb. 2018, doi: 10.1109/TSP.2018.2807417.
- [86] L. Wielandner, A. Venus, T. Wilding, and E. Leitinger, "Multipath-based SLAM with multiple-measurement data association," in *2023 26th Int. Conf. Inf. Fusion*, Jun. 27–30, 2023, doi: 10.23919/FUSION52260.2023.10224195.
- [87] M. G. Christensen and A. Jakobsson, *Multi-Pitch Estimation*, ser. Synthesis Lectures on Speech & Audio Processing, B. H. Juang, Ed. San Rafael, CA, USA: Morgan & Claypool, 2009.
- [88] Z. Duan and B. Pardo. Bach 10 dataset. Accessed: Dec. 2022. [Online]. Available: www.ece.rochester.edu/projects/air/datasets/Bach10.zip
- [89] A. de Cheveigné and H. Kawahara, "YIN, a fundamental frequency estimator for speech and music," *J. Acoust. Soc. Amer.*, Apr. 2022, doi: 10.1121/1.458024.
- [90] E. Benetos and T. Weyde, "An efficient temporally-constrained probabilistic model for multiple-instrument music transcription," in *Proc. 16th Int. Soc. Music Inf. Retrieval Conf.*, Malaga, Spain, Oct. 26–30, 2015.
- [91] F. Elvander, J. Swärd, and A. Jakobsson, "Online estimation of multiple harmonic signals," *IEEE/ACM Trans. Audio, Speech, Language Process.*, Feb. 2017, doi: 10.1109/TASLP.2016.2634118.
- [92] T. L. Hansen, P. B. Jørgensen, M.-A. Badiu, and B. H. Fleury, "An iterative receiver for OFDM with sparsity-based parametric channel estimation," *IEEE Trans. Signal Process.*, Oct. 2018, doi: 10.1109/TSP.2018.2868314.
- [93] E. Riegler, G. E. Kirkelund, C. N. Manchon, M. A. Badiu, and B. H. Fleury, "Merging belief propagation and the mean field approximation: A free energy approach," *IEEE Trans. Inf. Theory*, Jan. 2013, doi: 10.1109/tit.2012.2218573.
- [94] R. Prasad, C. R. Murthy, and B. D. Rao, "Joint channel estimation and data detection in MIMO-OFDM systems: A sparse Bayesian learning approach," *IEEE Trans. Signal Process.*, Oct. 2015, doi: 10.1109/TSP.2015.2451071.

Part II

Included Papers



Car Occupancy Detection Using Ultra-Wideband Radar

Jakob Möderl, Franz Pernkopf, and Klaus Witrisal

*2021 18th European Radar Conference (EuRAD), London, U.K., Apr. 5–7, 2022,
doi: 10.23919/EuRAD50154.2022.9784505*

Abstract—In this paper, we propose car occupancy detection using ultra-wideband (UWB) radar to detect the breathing motion of a person. The target signal and multipath propagation inside the car are modeled as a pinhole channel. We show that the received signal can be described by the Kronecker product of the channel and breathing motion vectors. Thus, the covariance can be computed from the delay profile of the channel and the power spectrum of the chest motion. An estimator-correlator is applied for the detection. Measurements have been performed to confirm the structural assumptions about the received signal whereas Monte-Carlo simulations are used to evaluate the detection performance. The simulations demonstrate an improved detection rate in low signal-to-noise ratio (SNR) conditions compared to a windowed energy detector or FFT-detector.

This research was partly funded by the Austrian Research Promotion Agency (FFG) within the project SEAMAL Front (project number: 880598).




Figure A.1: Pinhole channel model for the target signal.

A.1 Introduction

New safety regulations will require future cars to be able to detect if they are occupied when being locked in order to prevent the confinement of small children [1]. We propose to use the ultra-wideband (UWB) nodes of the keyless entry system for this use case as radar sensors. Recent works applied machine learning for this task, e.g. [2] and [3]. However, a disadvantage of such algorithms is that they are limited to the scenarios present in the training data. Thus, such models might have difficulties if they are deployed to other cars.

To overcome this limitation, we model the radio channel of the car and develop a detection algorithm based on this model. Hence, the algorithm can be deployed to any car model with little to no adaptation.

Notation: We use upper-case bold letters \mathbf{A} to denote matrices and lower-case bold letters \mathbf{a} to denote column vectors. $[\mathbf{A}]_{mn}$ denotes the (m, n) -th element of the matrix \mathbf{A} . \mathbf{A}^T denotes the transpose of \mathbf{A} , and \mathbf{A}^H denotes the Hermitian transpose of \mathbf{A} . $\mathbb{E}\{a\}$ denotes the expectation of the random variable a whereas a^* is used to denote the complex conjugate of a . The matrix \mathbf{I}_k is the $k \times k$ identity matrix. The convolution of two functions $a(\tau)$ and $b(\tau)$ is denoted as $(a * b)(\tau)$.

A.2 Signal Model

We consider a raised-cosine baseband UWB pulse $s(\tau)$ with pulse width T_p , which is transmitted at multiple time instances $t = mT_m$ at a carrier frequency f_c . This pulse propagates through the passenger cabin of the car where it is reflected from the windows, seats and other objects within the car. A receive antenna captures the reflected signal. If the car is occupied, the occupant reflects the signal as well and imposes some time variations on the reflected signal. The car is assumed to be parked, therefore the environment inside as well as outside of the car are assumed constant during the measurement period. Thus, we argue, the only change in the propagation environment is due to the presence of human beings.

A.2.1 Target Channel Model

Assuming a single point target, the received signal from such a target can be modeled using a pinhole channel as shown in Fig. A.1. In this approach the channel consists of three separate parts: A forward channel $h_f(\tau)$, the target (or pinhole) channel $h_p(t, \tau)$ and the backwards channel $h_b(\tau)$.

The forward channel $h_f(\tau) = \delta(\tau - \tau_d) + \nu(\tau)$ describes the propagation from the antenna to the target and is modeled by a deterministic line-of-sight component

with propagation delay τ_d and a stochastic part $\nu(\tau)$. The latter describes the diffuse cluster of reflections which arises from the signal being reflected of the vehicle interior. It is modeled as a complex Gaussian random process, which is characterized by its autocorrelation function

$$\mathbb{E}\{\nu(\tau)\nu^*(u)\} = S_\nu(\tau - \tau_d)\delta(\tau - u). \quad (\text{A.1})$$

Here $S_\nu(\tau)$ is the power delay profile (PDP) of the diffuse multipath (DM) components. For $\tau > 0$, it is modeled as an exponential decay [4–7]

$$S_\nu(\tau) = \frac{1}{K_{LoS} \cdot \gamma_0} e^{-\frac{\tau}{\gamma_0}}. \quad (\text{A.2})$$

The parameters of the PDP are the ratio $K_{LoS} = \frac{E_{LoS}}{E_{DM}}$ between energy of the line-of-sight component, which is normalized to $E_{LoS} = 1$, and the energy of the DM $E_{DM} = \int_0^\infty S_\nu(\tau) d\tau = \frac{1}{K_{LoS}}$, as well as the decay constant γ_0 .

The backward channel $h_b(\tau)$ is modeled in a similar manner as the forward channel. Due the commutativity of the convolution, we can summarize the forward and backward channels in a compound channel $h_{fb}(\tau) = (h_f * h_b)(\tau)$. The interested reader is referred to [8] for the calculation of the compound PDP $S_{fb}(\tau)$ and its parameters. This results in an distorted pulse $s_h(\tau) = (s * h_{fb})(\tau)$ that interacts with the target. In case the movement of the target is smaller than one half of the wavelength λ , the channels $h_f(\tau)$ and $h_b(\tau)$ can be assumed constant over the course of the measurement [9].

A.2.2 Target Model

The occupant is modeled as a single point target which reflects the signal with a certain reflection coefficient $\alpha \in \mathbb{R}$. Furthermore, the distance between the target and the antenna changes slightly over time due to the small movements $b(t)$ living beings inherently make, most predominantly the chest motion due to breathing. Thus, the signal reflected by the target is shifted by an additional delay $\tau_b(t) = 2\frac{b(t)}{c}$ resulting in the occupant channel

$$h_p(t, \tau) = \alpha e^{-j2\pi f_c \tau_b(t)} \delta(\tau - \tau_b(t)). \quad (\text{A.3})$$

We model $b(t)$ as a wide-sense stationary, zero-mean Gaussian process. Its autocorrelation $R_{bb}(t') = \mathbb{E}\{b(t+t')b(t)\} = \int_{-\infty}^\infty S_{bb}(f)e^{j2\pi t'f} df$ is defined by its corresponding power spectral density $S_{bb}(f)$. Even though $S_{bb}(f)$ is unknown in practice, we approximate it for the purpose of the detection as a double-sided rectangular spectrum $S_{bb}(f) = \frac{B^2}{2(f_2 - f_1)}$ for $f_1 \leq |f| \leq f_2$ and 0 elsewhere. Here B^2 is the average power of $b(t)$ and f_1, f_2 define the “band of interest” of the detector and should be chosen such that the breathing rate of babies as well as adults are within the selected bandwidth. It follows, that the received target signal $r_p(t, \tau)$ is given by the convolution of the distorted pulse $s_h(\tau)$ with the target channel $h_p(\tau)$ resulting in

$$r_p(t, \tau) = \alpha e^{-j4\pi \frac{b(t)}{\lambda_c}} s_h(\tau - \tau_b(t)). \quad (\text{A.4})$$

In this equation $\lambda_c = \frac{c}{f_c}$ denotes the carrier wavelength.

This can be understood as the distorted pulse $s_h(\tau)$ being shifted in time and phase modulated due to the body movement $b(t)$. For small body movements, such that $4\pi \frac{b(t)}{\lambda_c} \ll 1$ and $\tau_b(t) \ll T_p$, this can be approximated with

$$e^{-j4\pi \frac{b(t)}{\lambda_c}} s_h(\tau - \tau_b(t)) \approx \left(1 - j4\pi \frac{b(t)}{\lambda_c}\right) s_h(\tau), \quad (\text{A.5})$$

resulting in the received target signal

$$r_p(t, \tau) = \beta b(t) s_h(\tau) + \alpha s_h(\tau) \quad (\text{A.6})$$

denoting $\beta = -j4\pi \frac{b(t)}{\lambda_c}$. The received target signal consist of a static part $\alpha s_h(\tau)$ that does not depend on the time t and will be removed during clutter removal, as well as a dynamic part $y(t, \tau) = \beta b(t) s_h(\tau)$, which we will focus on for the detection.

A.2.3 Clutter and Noise

Noise $w(t, \tau)$ and clutter $c(\tau)$ will be received in addition to the signal from the target $r_p(t, \tau)$. The noise $w(t, \tau)$ is modeled as a circularly-symmetric complex white Gaussian noise with power N_0 . The clutter $c(\tau)$ models all propagation paths that do not interact with the target. Furthermore, since the clutter stems from the same environment as $h_f(\tau)$ and $h_b(\tau)$ it can be modeled in a similar fashion. However, due to the clutter being removed in the pre-processing, no further effort was put into modeling the exact shape of the clutter.

Finally, the complete received signal is given by

$$r(t, \tau) = y(t, \tau) + \alpha s_h(\tau) + c(\tau) + w(t, \tau). \quad (\text{A.7})$$

A.2.4 Clutter Removal and Detection

In order to subtract the clutter, the mean $\bar{r}(\tau)$ is calculated over slow time and subtracted from the received signal $x(t, \tau) = r(t, \tau) - \bar{r}(\tau)$. The signal is sampled at several instances in fast time $\tau = nT_s$, $n \in [1, 2, \dots, N]$ and slow time $t = mT_m$, $m \in [1, 2, \dots, M]$, given the sampling intervals T_m and T_s . The samples of $x(mT_m, nT_s)$ are stacked into a large vector $\mathbf{x} = [\mathbf{x}_1^T, \dots, \mathbf{x}_N^T]^T$ composed from the subvectors $\mathbf{x}_i = [x(1, iT_s), \dots, x(MT_m, iT_s)]^T$ of the received clutter-free signal at fast time $\tau = iT_s$. Let \mathbf{y} and \mathbf{w} be vectors that are constructed from $y(t, \tau)$ and $w(t, \tau)$, respectively, in the same manner as \mathbf{x} is constructed from $x(t, \tau)$, such that $\mathbf{x} = \mathbf{y} + \mathbf{w}$.

Furthermore, let $\mathbf{b} = \beta \cdot [b(T_m), b(2T_m), \dots, b(MT_m)]^T$ be a column vector of the sampled breathing signal and $\mathbf{h} = [s_h(T_s), s_h(2T_s), \dots, s_h(NT_s)]^T$ a column vector of the sampled distorted pulse, then it follows that $\mathbf{y} = \mathbf{h} \otimes \mathbf{b}$ is the Kronecker product of these vectors with covariance $\mathbf{C}_{yy} = \mathbb{E}\{\mathbf{y}\mathbf{y}^H\} = \mathbf{C}_{hh} \otimes \mathbf{C}_{bb}$. If the samples of \mathbf{y} are stacked in a matrix $\mathbf{Y} = \mathbf{b}\mathbf{h}^T$ instead of a large vector, this matrix is the outer product of the breathing and channel vectors \mathbf{b} and \mathbf{h} and, therefore, it is of rank 1.

The covariance $\mathbf{C}_{hh} = \mathbb{E}\{\mathbf{h}\mathbf{h}^H\}$ can be calculated using $\mathbf{C}_{hh} = \mathbf{S}\mathbf{S}_{fb}\mathbf{S}^H$, where \mathbf{S}_{fb} is a diagonal matrix which is composed of the sampled PDP $S_{fb}(nT_s)$ on its main diagonal and $\mathbf{S} = [\mathbf{s}_0, \mathbf{s}_1, \dots, \mathbf{s}_{N-1}]$ is a convolution matrix of the sampled

pulses $\mathbf{s}_i = [s(iT_s), s((i+1)T_s), \dots, s((i+N-1)T_s)]^T$. Note, that \mathbf{C}_{hh} depends implicitly on the distance to the target τ_d which we skipped in the notation to be more concise. The elements of the covariance matrix $\mathbf{C}_{bb} = \mathbb{E}\{\mathbf{b}\mathbf{b}^H\}$ can be calculated as $[\mathbf{C}_{bb}]_{ij} = |\beta|^2 R_{bb}((i-j)T_m)$.

Since \mathbf{y} is a random vector with known covariance \mathbf{C}_{yy} , we face the following detection problem

$$\begin{cases} \mathcal{H}_0 : \mathbf{x} = \mathbf{w} & \sim \mathcal{CN}(\mathbf{0}, \sigma^2 \mathbf{I}_{NM}) \\ \mathcal{H}_1 : \mathbf{x} = \mathbf{y} + \mathbf{w} & \sim \mathcal{CN}(\mathbf{0}, \mathbf{C}_{yy} + \sigma^2 \mathbf{I}_{NM}). \end{cases}$$

In this equation \mathcal{H}_0 denotes the case of no target occupying the car and \mathcal{H}_1 denotes the case where there is a person in the car. The optimal decision criterion for this kind of detection problem is given by the estimator-correlator which decides for \mathcal{H}_1 if the following test statistics $T(\mathbf{x})$ is greater than a threshold γ [10]

$$T(\mathbf{x}) = \mathbf{x}^H \mathbf{C}_{yy} (\mathbf{C}_{yy} + \sigma^2 \mathbf{I}_{NM})^{-1} \mathbf{x}. \quad (\text{A.8})$$

Since the distance to the target τ_d is unknown, the received signal is successively shifted in fast time while $T(\mathbf{x})$ is calculated in each step. The maximum over all $T(\mathbf{x})$ is then compared to the threshold γ .

A.3 Radar Measurements

In order to confirm our assumptions, a radar measurement has been performed inside a sedan (Lancia Thema 2011) to characterize the received signal \mathbf{x} . The channel impulse responses (CIRs) have been measured using a correlative channel sounder [11] utilizing a raised cosine pulse $s(\tau)$ with a center frequency of $f_c = 7$ GHz and a pulse width of $T_p = 0.5$ ns (2 GHz bandwidth). The measurement equipment including channel sounder, cables and connectors has been calibrated, but not the antennas. Dipole antennas with circular patches have been used [12, Fig. B.5.(b)]. The receive and transmit antennas have been placed next to each other on the ceiling of the car between the front-row seats. The antennas are angled such that the nulls in z-direction of their antenna patterns are pointed towards each other in an effort to reduce the coupling between the antennas. An adult subject was sitting in the back seat at a distance of 1 m to the antennas and has been instructed to sit natural, keep the arms at the side of their torso and to breathe normally for the measurement duration of 10 s. During the measurement a CIR was captured every $T_m = 0.1$ s.

The received signal after clutter removal is shown in Fig. A.2 with the breathing activity clearly visible at a fast time of $\tau \approx 7$ ns. The effect of the channel “spreads” the breathing activity in fast time after the line-of-sight component. It can be observed, that the columns in Fig. A.2 are approximately the same over different values of the fast time τ , except for their amplitude. An analysis of the eigenvalues of \mathbf{X} revealed that, after accounting for the noise energy, approximately 87 % of the signal energy is concentrated in the first eigenvalue and, thus, by approximating \mathbf{X} with a matrix of rank 1. This indicates that our modeling assumptions are correct.




Figure A.2: Clutter-free received signal arranged as a matrix $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N]$ of a radar measurement of an adult sitting inactive on the backseat of a vehicle. The respiratory motion is visible at a fast time of $\tau \approx 7$ ns.

A.4 Performance Characterization

In order to characterize the performance of the detector, a Monte-Carlo simulation with $N_{mc} = 10^5$ runs has been performed. The body motion was approximated as a sinusoidal function $b(t) = A \cos(2\pi f_r t + \varphi_r)$. The respiration frequency f_r is drawn uniformly from 20 – 40 breaths per minute (bpm) and the phase drawn uniformly from $\varphi_r \in [0, 2\pi)$. The amplitude was fixed to $A = 1$ mm. This corresponds to the approximate breathing rate and amplitude of an infant. Since we want to detect the breathing motion of adults, which typically ranges from 10 – 20 bpm ($\frac{1}{6} - \frac{2}{6}$ Hz), as well as babies, which might range as high as 60 bpm (1 Hz) we use $f_1 = \frac{1}{6}$ Hz and $f_2 = 1$ Hz as bandwidth for the detector.

The power P_R received from a radar target at distance d is described by the radar range equation

$$P_R = P_T G_T \frac{\sigma_{RCS}}{4\pi d^2} \frac{A_{eff}}{4\pi d^2}. \quad (\text{A.9})$$

In this equation P_T is the transmit power, G_T the gain of the transmit antenna, A_{eff} the effective area of the receive antenna and σ_{RCS} is the radar cross section of the target. Using (A.9), the reflection coefficient α can be calculated as

$$\alpha = \sqrt{\frac{P_R}{P_T}} = \frac{\sqrt{G_T \cdot A_{eff} \cdot \sigma_{RCS}}}{4\pi d^2}. \quad (\text{A.10})$$

Note, that the radar cross section σ_{RCS} refers only to the parts of the targets body that move, i.e. the chest area, instead of the radar cross section of the whole body. A chest area of 15×10 cm is assumed for an infant. Assuming isotropic antennas we have $G_T = 1$. The effective area of the antenna $A_{eff} = 4 \text{ cm}^2$ is chosen to be similar to the antennas used in the measurements. The parameters $K_{LoS} = 0.75$ and $\gamma_0 = 20$ ns of the forward and backward channel as well as the distance between the antenna and the target $d = 1$ m are chosen in accordance with the measurements.

In addition to the developed “rank-1 signal” detector, the performance of a windowed energy detector as suggested by [13] and an FFT-detector are evaluated




Figure A.3: Detection probability for the developed rank-1 signal detector compared to a windowed energy detector and FFT-detector.

for comparison. The windowed energy detector computes the energy over a window in fast time with length L . Due to the unknown distance to the target, the window is shifted in fast time and the maximum compared against a threshold. This detector does not make any assumptions about the received signal, except for it being concentrated around a window of delays. In the simulations a window of length $L = 46$ samples is used, which corresponds to a duration of 6.67 ns.

The FFT-Detector computes the discrete Fourier transform over slow time and compares the peak value against a threshold. This type of detector is optimal for periodic signals [10], such as the breathing motion in the simulation. However, this detector does not incorporate knowledge about the channel into the detection. Furthermore, the breathing motion will rarely be strictly periodic [14] which will decrease the performance of the FFT-Detector in practice.

A.5 Results

To assess the performance of the detectors, first a simulation was performed without any target present to calculate the thresholds such that a false alarm probability of $P_{FA} = 0.01$ is achieved by all detectors. Next, simulations including a target have been performed, during which the reflection coefficient α was varied such that the receive signal-to-noise ratio (SNR) $\frac{E_s}{N_0} |\beta A|^2$ ranges from -6 to 12 dB. Here $E_s = \int_{-\infty}^{\infty} s^2(\tau) d\tau$ is the energy of the pulse $s(\tau)$. The resulting detection probabilities are depicted in Fig. A.3.

The detection probability for the windowed energy detector is the worst, which is as expected since it does not make any assumptions about the signal except for it being concentrated in a window of delays. The FFT-detector makes assumptions about the form of the signal in slow time (i.e. a periodic breathing movement) but not in fast time and, thus, it is able to improve the transition region by approximately 3 dB compared to the windowed energy detector. The rank-1 signal detector

incorporates knowledge about the slow time as well as fast time of the received signal and is therefore able to outperform the other two, resulting in an improvement of the transition region of 6 dB compared to the windowed energy detector and 3 dB compared to the FFT-Detector, respectively.

A.6 Discussion

The results show good detection probabilities for a target that is not moving except for the breathing motion. By modeling the breathing movement as a random signal, the performance of the rank-1 signal detector should not be altered by arrhythmic breathing in contrast to the FFT-detector which relies on a periodic signal. Due to the simple structure of the detector and the target signal being modeled as rank-1 matrix, it should be possible to apply this detector to any car model without further adaptation.

However, if the assumption of an inactive target does not hold, e.g. if the target is moving its limbs, the detector might not perform as expected. In practice this might be compensated by the larger amplitude of limb movement compared to the breathing motion. Extending the developed detector to cope with a generic moving target is part of ongoing research.

Similarly, if more than one target is present, the rank-1 assumptions is also not fulfilled. However, it is unlikely that the signals from both targets interfere destructively during the whole measurement in both fast time and slow time domain. Thus, we argue, that in this case the performance of the rank-1 detector will be comparable to the performance that was shown for a single occupant.

Since UWB signals show good penetration properties for “light” materials, the performance of the detector will not be impacted if the occupant is covered e.g. by a blanket. However, if the chest of the occupant is not facing towards the radar sensor, the radar cross section of the moving area is reduced which will result in a reduced SNR and decreased detection probability.

Acknowledgment

This research was partly funded by the Austrian Research Promotion Agency (FFG) within the project SEAMAL Front (project number: 880598).

Bibliography

- [1] Euro NCAP. (2021, May) Euro NCAP 2025 roadmap. [Online]. Available: <https://cdn.euroncap.com/media/30701/euroncap-roadmap-2025-v4-print.pdf>
- [2] S. Lim, S. Lee, J. Jung, and S.-C. Kim, “Detection and localization of people inside vehicle using impulse radio ultra-wideband radar sensor,” *IEEE Sensors J.*, Apr. 2020, doi: 10.1109/JSEN.2019.2961107.
- [3] Y. Ma, Y. Zeng, and V. Jain, “CarOSense: Car occupancy sensing with the ultra-wideband keyless infrastructure,” *Proc. ACM Interact. Mobile Wearable Ubiquitous Technol.*, Sep. 2020, doi: 10.1145/3411820.

- [4] O. Delangre, S. Van Roy, P. De Doncker, M. Lienard, and P. Degauque, "Modeling in-vehicle wideband wireless channels using reverberation chamber theory," in *IEEE 66th Veh. Technol. Conf.*, Sep. 30–Oct. 3, 2007, doi: 10.1109/VETEFCF.2007.451.
- [5] F. Bellens, F. Quitin, F. Horlin, and P. De Doncker, "UWB channel analysis within a moving car," in *9th Int. Conf. Intell. Transp. Syst. Telecommun.*, Oct. 20–22, 2009, doi: 10.1109/ITST.2009.5399271.
- [6] Y. Katayama, K. Terasaka, K. Higashikaturagi, I. Matunami, and A. Kajiwara, "Ultra-wideband impulse-radio propagation for in-vehicle wireless link," in *IEEE 64th Veh. Technol. Conf.*, Sep. 25–28, 2006, doi: 10.1109/VTCF.2006.63.
- [7] T. Kobayashi, "Measurements and characterization of ultra wideband propagation channels in a passenger-car compartment," in *IEEE 9th Int. Symp. Spread Spectr. Techn. and Appl.*, Aug. 28–31, 2006, doi: 10.1109/ISSSTA.2006.311768.
- [8] D. Arnitz, U. Muehlmann, and K. Witrisal, "Wideband characterization of backscatter channels: Derivations and theoretical background," *IEEE Trans. Antennas Propag.*, Jan. 2012, doi: 10.1109/TAP.2011.2167923.
- [9] N. Michelusi, U. Mitra, A. F. Molisch, and M. Zorzi, "UWB sparse/diffuse channels, part I: Channel models and bayesian estimators," *IEEE Trans. Signal Process.*, Jun. 2012, doi: 10.1109/TSP.2012.2205681.
- [10] S. M. Kay, *Fundamentals of Statistical Signal Processing: Detection Theory*, 14th ed., A. V. Oppenheim, Ed. Upper Saddle River, NJ, USA: Prentice-Hall, 2009.
- [11] ILMSENS. (2021, May) m:explore - datasheet. [Online]. Available: <https://www.ilmsens.com/app/download/14567836924/SH-3100%20-%20mexplore%20datasheet%2020170728.pdf>
- [12] C. Krall, "Signal processing for ultra wideband transceivers," phdthesis, Signal Process. and Speech Commun. Lab., Graz Univ. of Technol., Graz, Austria, 2008.
- [13] Y. Kilic, H. Wymeersch, A. Meijerink, M. J. Bentum, and W. G. Scanlon, "Device-free person detection and ranging in UWB networks," *IEEE J. Sel. Topics Signal Process.*, Feb. 2014, doi: 10.1109/JSTSP.2013.2281780.
- [14] B. Schleicher, I. Nasr, A. Trasser, and H. Schumacher, "IR-UWB radar demonstrator for ultra-fine movement detection and vital-sign monitoring," *IEEE Trans. Microw. Theory Techn.*, May 2013, doi: 10.1109/TMTT.2013.2252185.



Variational Message Passing-based Respiratory Motion Estimation and Detection Using Radar Signals

Jakob Möderl, Erik Leitinger, Franz Pernkopf, and Klaus Witrisal

2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Rhodes, Greece, Jun. 4–10, 2023, doi: ICASSP49357.2023.10095725

Abstract—We present a variational message passing (VMP)-based approach to detect the presence of a person based on their respiratory chest motion using multistatic ultra-wideband (UWB) radar. In the process, the respiratory motion is estimated for contact-free vital sign monitoring. The received signal is modeled as a backscatter channel and the respiratory motion and propagation channels are estimated using VMP. We use the evidence lower bound (ELBO) to approximate the model evidence for the detection. Numerical analyses and measurements demonstrate that the proposed method leads to a significant improvement in the detection performance compared to a fast Fourier transform (FFT)-based detector or an estimator-correlator in low-signal-to-noise ratio (SNR) conditions, since the multipath components (MPCs) are better incorporated into the detection procedure. Specifically, the proposed method has a detection probability of 0.95 at -20 dB SNR, while the estimator-correlator and FFT-based detector have 0.32 and 0.05, respectively.

This research was partly funded by the Austrian Research Promotion Agency (FFG) within the project SEAMAL Front (project number: 880598).

B.1 Introduction

Future cars will be required to detect if they are occupied when the car is being locked to prevent the confinement of small children, which makes detecting the presence of people a safety-critical application [1–3]. The respiratory motion of the chest provides a dynamic feature used to separate the target’s radar response from the strong clutter present in this use case [4]. We derive a Kronecker-factorized signal model for vital sign estimation in (strong) clutter using ultra-wideband (UWB) radar and present a variational message passing (VMP) algorithm [5] [6, Ch. 10] to detect the presence of people and estimate the respiratory chest motion. An advantage of the derived model is, that it is linear in both the respiratory motion as well as the propagation channel and, thus, the message passing equations can be solved analytically.

Recent approaches to detect and estimate the respiratory chest motion based on radar responses are mostly based on the intuition that the respiratory motion is periodic. They apply techniques such as fast Fourier transform (FFT) processing, principal component analysis or energy detection [7–11]. However, this assumption is easily violated, e.g. by infants who regularly experience apnea (short pauses with no respiration) [12]. This is especially important, considering that infants are very hard to detect in the first place, due to the small radar cross section and respiratory motion amplitude. On the other hand, data-driven approaches [13, 14] are typically limited by the small and heterogeneous data sets available. These works do not explicitly model the propagation of multipath components (MPCs). However, in a tightly enclosed space, such as the interior of a car, the MPCs that interact with the target carry additional information and can be used to increase the signal-to-noise ratio (SNR), which is of critical importance in the given use case. In this work, we apply VMP to improve upon the results of [4], which already incorporates MPCs into the detection.

Notation: We define \mathbf{I}_n to be the $n \times n$ identity matrix and $\mathbf{1} = [1, 1, \dots, 1]^T$ to be a vector of ones with appropriate length. We use \odot and \otimes to denote the Hadamard (element wise) and Kronecker product of two vectors or matrices, respectively. The real operator and matrix trace operator are denoted as $\text{Re}\{\cdot\}$ and $\text{tr}(\cdot)$. We use $\mathcal{N}(\mathbf{a}|\mathbf{b}, \mathbf{C})$ and $\mathcal{CN}(\mathbf{a}|\mathbf{b}, \mathbf{C})$ to denote that the vector \mathbf{a} is distributed with a (complex) multivariate Gaussian distribution with mean \mathbf{b} and covariance matrix \mathbf{C} . Similarly, $\text{Ga}(a|d, e)$ is used to denote that the variable a is gamma-distributed with shape parameter d and rate parameter e . The differential entropy of the distribution $q(x)$ is denoted as $\mathbb{H}(p)$ and the expectation of the function $f(x)$ with respect to $q(x)$ as $\langle f(x) \rangle_{q(x)}$.

B.2 Signal Model

We consider the case of a person sitting in a car, without intentional body movement. However, the chest of the person expands and contracts continuously due to the persons respiration. We propose to model the chest movement in direction of the antenna $b_t(t)$ as the realization of a zero-mean Gaussian random process. We aim to detect the presence of the person and estimate $b_t(t)$ using multistatic UWB radar signals.

Let B be the bandwidth of the transmit pulse and \mathbf{s} be a vector of N samples, equally spaced with spacing $\Delta f = B/N$, of the complex baseband representation of the transmit pulse $s(f)$ centred at carrier frequency f_c . Several repetitions of the pulse are transmitted at times $t = mT_{\text{rep}}$, $m \in \{0, 1, \dots, M-1\}$ from the transmit antenna. After propagating over K time-varying channels with frequency response $\mathbf{h}_k(t)$, each signal $\mathbf{r}_{k,m} = \mathbf{h}_k(mT_{\text{rep}}) \odot \mathbf{s} + \mathbf{w}_{k,m}$ received at antenna k at repetition m is corrupted by noise $\mathbf{w}_{k,m}$. The noise samples $\mathbf{w}_{k,m}$ are generated by a noise process $W_{m,k}(f)$, which is modeled as additive proper complex white Gaussian noise with double-sided power spectral density $N_0/2$, and is assumed to be independent across m , k and f . Thus, $\mathbf{w}_{k,m}$ is a circular symmetric proper complex Gaussian random vector with covariance $\mathbf{C}_w = \lambda^{-1} \mathbf{I}_N$ and precision $\lambda = T_s/N_0$.

In order to remove the clutter, the signals from all antennas at time m are stacked $\mathbf{r}_m = [\mathbf{r}_{1,m}^T, \mathbf{r}_{2,m}^T, \dots, \mathbf{r}_{K,m}^T]^T$ and the mean $\bar{\mathbf{r}} = \frac{1}{M} \sum_{m=0}^{M-1} \mathbf{r}_m$ over m is subtracted $\tilde{\mathbf{r}}_m = \mathbf{r}_m - \bar{\mathbf{r}}$. Finally, the signals are stacked into a large column vector $\tilde{\mathbf{r}} = [\tilde{\mathbf{r}}_0^T, \tilde{\mathbf{r}}_1^T, \dots, \tilde{\mathbf{r}}_{M-1}^T]^T$. The time-varying part of the received signal

$$\tilde{\mathbf{r}} = \mathbf{b}_t \otimes \mathbf{h}_s + \mathbf{w} \quad (\text{B.1})$$

is given as the product of the respiratory motion $\mathbf{b}_t = [b_t(0), b_t(T_{\text{rep}}), \dots, b_t((M-1)T_{\text{rep}})]^T \in \mathbb{R}^M$ and a stacked channel vector $\mathbf{h}_s = [\mathbf{h}_{s,1}^T, \mathbf{h}_{s,2}^T, \dots, \mathbf{h}_{s,K}^T]^T$ in additive white Gaussian noise $\mathbf{w} = [\mathbf{w}_{1,0}^T, \mathbf{w}_{2,0}^T, \dots, \mathbf{w}_{K,M-1}^T]^T$. We provide a detailed derivation of (B.1) in Section B.2.1.

Since we apply a frequency selective prior to \mathbf{b}_t the resulting covariance $\mathbf{C}_{\mathbf{b}_t}$ is not full rank. Thus, all computations are performed in the eigenspace $\mathbf{b} = \mathbf{U}^T \mathbf{b}_t$ corresponding to the eigendecomposition $\mathbf{C}_{\mathbf{b}_t} = \mathbf{U} \mathbf{C}_b \mathbf{U}^T$, where \mathbf{C}_b is a diagonal matrix with all L non-zero eigenvalues of $\mathbf{C}_{\mathbf{b}_t}$ on its main diagonal and \mathbf{U} is a matrix whose columns contain the corresponding eigenvectors. Let $\mathbf{e}_k = [0, \dots, 0, 1, 0, \dots, 0]^T$ be a vector of length K with all zeros except for an 1 at the k -th position and $\mathbf{w}_k = [\mathbf{w}_{k,0}^T, \mathbf{w}_{k,1}^T, \dots, \mathbf{w}_{k,M-1}^T]^T$, $\tilde{\mathbf{r}}$ can be expressed either as a linear function of the breathing signal $\mathbf{b}_t = \mathbf{U} \mathbf{b}$ given the block-diagonal matrix $\mathbf{H} = \mathbf{I}_M \otimes \mathbf{h}_s$ and the signal received at each antenna k after clutter removal $\tilde{\mathbf{r}}_{Ak} = ((\mathbf{e}_k^T \otimes \mathbf{I}_N) \otimes \mathbf{I}_M) \tilde{\mathbf{r}}$ as a linear function of $\mathbf{h}_{s,k}$ given the block-diagonal matrix $\mathbf{B} = \mathbf{U} \mathbf{b} \otimes \mathbf{I}_N$:

$$\tilde{\mathbf{r}} = \mathbf{H} \mathbf{U} \mathbf{b} + \mathbf{w} \quad (\text{B.2})$$

$$\tilde{\mathbf{r}}_{Ak} = \mathbf{B} \mathbf{h}_{s,k} + \mathbf{w}_k. \quad (\text{B.3})$$

Assuming independent channels $\mathbf{h}_{s,k}$, the likelihood of receiving $\tilde{\mathbf{r}}$ is $p(\tilde{\mathbf{r}} | \mathbf{b}, \mathbf{h}_s, \lambda) = \mathcal{CN}(\tilde{\mathbf{r}} | \mathbf{H} \mathbf{U} \mathbf{b}, \lambda^{-1} \mathbf{I}_{KNM}) = \prod_{k=1}^K \mathcal{CN}(\tilde{\mathbf{r}}_{Ak} | \mathbf{B} \mathbf{h}_{s,k}, \lambda^{-1} \mathbf{I}_{NM})$.

B.2.1 Propagation Environment and Target Model

Let $\mathbf{h}_{\text{fb},k}(t) = [h_{\text{fb},k}(-\frac{N}{2} \Delta f, t), h_{\text{fb},k}((-\frac{N}{2} + 1) \Delta f, t), \dots, h_{\text{fb},k}((\frac{N}{2} - 1) \Delta f, t)]^T$ be a vector of the sampled baseband frequency response $h_{\text{fb},k}(f, t)$ of the channel to receive antenna k at time t [15]. We model the propagation environment as a time-varying backscatter channel [16]

$$\mathbf{h}_k(t) = \mathbf{h}_{t,k}(t) \odot \mathbf{h}_{\text{fb},k} + \mathbf{h}_{c,k}. \quad (\text{B.4})$$





Figure B.1: Receive signal $\tilde{\mathbf{R}}$ of a UWB radar measurement of an adult sitting on the back seat of a car.

Introducing briefly the target channel $\mathbf{h}_{t,k}(t)$, the MPCs which interact with the target are modeled by a forward-backward channel $\mathbf{h}_{fb,k} = \mathbf{h}_{f,k} \odot \mathbf{h}_{b,k}$, which is the product of a forward channel $\mathbf{h}_{f,k}$ covering the propagation from the transmit antenna to the target and a backward channel $\mathbf{h}_{b,k}$ covering the propagation from the target back to the receive antenna. All other received MPCs are termed as clutter and are modeled by the frequency response $\mathbf{h}_{c,k}$. Since the target is assumed to be stationary, the channels $\mathbf{h}_{f,k}$, $\mathbf{h}_{b,k}$ and $\mathbf{h}_{c,k}$ can be assumed time-invariant as long as the respiratory motion $b_t(t)$ is much smaller than the smallest wavelength of the transmit signal.

The target is modeled as a single point target with a time-varying baseband frequency response [15] $h_{t,k}(f, t) = \alpha_k e^{-j2\pi(f+f_c)\tau_{b,k}(t)}$, representing the reflection of the incoming signal by a coefficient $\alpha_k \in \mathbb{C}$ and a time-varying delay $\tau_{b,k}(t)$. Let c be the propagation speed of the signal and $0 < \rho_k < 2$ a geometry dependent coefficient such that $\tau_{b,k}(t) = \rho_k b_t(t)/c$. Assuming $b_t(t) \ll \frac{c}{2\pi\rho_k(f+f_c)}$, we apply a first order Taylor-approximation $e^{-j2\pi(f+f_c)\tau_{b,k}(t)} \approx 1 - j2\pi(f+f_c)\tau_{b,k}(t)$. Let $\mathbf{f} = [-\frac{N}{2}\Delta f, (-\frac{N}{2}+1)\Delta f, \dots, (\frac{N}{2}-1)\Delta f]^T$, the sampled frequency response of the target changes over time as $\mathbf{h}_{t,k}(t) = \alpha_k(\mathbf{1} - j2\pi\rho_k\frac{\mathbf{f}+f_c\mathbf{1}}{c}b_t(t))$. Let $\mathbf{h}_{s,k} = -j2\pi\rho_k\alpha_k/c \cdot \mathbf{h}_{fb,k} \odot (\mathbf{f} + f_c\mathbf{1}) \odot \mathbf{s}$, the signal received at antenna k at time m is $\mathbf{r}_{m,k} = b_t(mT_{\text{rep}}) \cdot \mathbf{h}_{s,k} + \mathbf{w}_{m,k} + \text{const}$. Thus, after removing the constant term we arrive at (B.1).

Fig. B.1 shows a radar measurement from the experiments described in Section B.4. The receive signal is stacked into a matrix $\tilde{\mathbf{R}} = [\tilde{\mathbf{r}}_0, \tilde{\mathbf{r}}_1, \dots, \tilde{\mathbf{r}}_{M-1}]$. To highlight the MPCs, the channel is transformed to time domain by the inverse-FFT matrix \mathbf{V} , where the $[k, n]$ -th element of \mathbf{V} is defined as $\frac{1}{\sqrt{N}}e^{j2\pi nk/N}$. Each row is approximately a scaled version of the row containing the line-of-sight (LoS), as predicted by the derived signal model.

Figure B.2: Factor graph of the posterior distribution $p(\mathbf{b}, \mathbf{h}_s, \lambda | \tilde{\mathbf{r}})$.

B.3 Variational Message Passing

Obtaining the maximum a-posteriori solutions for \mathbf{b} and \mathbf{h}_s is computationally infeasible due to the large dimensions of \mathbf{b} and \mathbf{h}_s . Therefore, we apply a structured mean-field approach to approximate the posterior distribution $p(\mathbf{b}, \mathbf{h}_s, \lambda | \tilde{\mathbf{r}}) \propto p(\tilde{\mathbf{r}} | \mathbf{b}, \mathbf{h}_s, \lambda) p(\mathbf{b}) p(\lambda) \prod_{k=1}^K p(\mathbf{h}_{s,k})$, which is illustrated in Fig. B.2, with a factorized distribution $q_1(\mathbf{b}, \mathbf{h}_s, \lambda) = q_b(\mathbf{b}) q_\lambda(\lambda) \prod_{k=1}^K q_{h,k}(\mathbf{h}_{s,k})$ [17]. VMP is applied to minimize the Kullback-Leibler divergence $\mathcal{D}_{\text{KL}}(q_1 || p(\mathbf{b}, \mathbf{h}_s, \lambda | \tilde{\mathbf{r}}))$ of the true posterior $p(\mathbf{b}, \mathbf{h}_s, \lambda | \tilde{\mathbf{r}})$ from $q_1(\mathbf{b}, \mathbf{h}_s, \lambda)$ by maximizing the evidence lower bound (ELBO) [18–21]. The ELBO is maximized using coordinate ascent, iteratively maximizing the ELBO with respect to one distribution $q_j \in \mathcal{Q} = \{q_b, q_\lambda, q_{h,1}, \dots, q_{h,K}\}$ by

$$q_j \propto \exp \left\{ \left\langle \ln p(\mathbf{b}, \mathbf{h}_s, \lambda | \tilde{\mathbf{r}}) \right\rangle_{q_j} \right\} \quad (\text{B.5})$$

while keeping the remaining distributions $q_{\bar{j}} = \prod_{q_k \in \mathcal{Q} \setminus q_j} q_k$ fixed. Note, that the fixed point can be found analytically, if conjugate priors are used. Therefore, we assume a zero-mean Gaussian prior $p(\mathbf{b}) = \mathcal{N}(\mathbf{b} | \mathbf{0}, \mathbf{C}_{b_0})$ and a complex proper zero-mean Gaussian prior $p(\mathbf{h}_{s,k}) = \mathcal{CN}(\mathbf{h}_{s,k} | \mathbf{0}, \mathbf{C}_{h_0,k})$ for the respiratory motion and channel vectors, respectively. Furthermore, we assume the scale-invariant Jeffrey's prior $p(\lambda) \propto \lambda^{-1}$ for the noise precision λ . The resulting distributions $q_b(\mathbf{b}) = \mathcal{N}(\mathbf{b} | \hat{\mathbf{b}}, \hat{\mathbf{C}}_b)$, $q_{h,k}(\mathbf{h}_{s,k}) = \mathcal{CN}(\mathbf{h}_{s,k} | \hat{\mathbf{h}}_{s,k}, \hat{\mathbf{C}}_{h,k})$, and $q_\lambda(\lambda) = \text{Ga}(\lambda | KNM, \hat{M}_\lambda)$ are fully described by the parameters $\hat{\mathbf{b}}, \hat{\mathbf{C}}_b, \hat{\mathbf{h}}_{s,k}, \hat{\mathbf{C}}_{h,k}$ and $\hat{M}_\lambda = KNM / \hat{\lambda}_1$. Let $\hat{E}_b^{[i]} = \text{tr}(\hat{\mathbf{C}}_b^{[i]}) + \|\hat{\mathbf{b}}^{[i]}\|^2$ and $\hat{E}_h^{[i]} = \sum_{k=1}^K \text{tr}(\hat{\mathbf{C}}_{h,k}^{[i]}) + \|\hat{\mathbf{h}}_{s,k}^{[i]}\|^2$, the following messages are computed at iteration i :

$$\hat{\mathbf{C}}_{h,k}^{[i]} = (\mathbf{C}_{h_0,k}^{-1} + \hat{\lambda}_1^{[i-1]} \hat{E}_b^{[i-1]} \mathbf{I}_N)^{-1} \quad (\text{B.6})$$

$$\hat{\mathbf{h}}_{s,k}^{[i]} = \hat{\lambda}_1^{[i-1]} \hat{\mathbf{C}}_{h,k}^{[i]} \hat{\mathbf{B}}^{[i-1] \text{H}} \tilde{\mathbf{r}}_{Ak} \quad (\text{B.7})$$

$$\hat{\mathbf{C}}_b^{[i]} = (\mathbf{C}_{b_0}^{-1} + 2\hat{\lambda}_1^{[i-1]} \hat{E}_h^{[i]} \mathbf{I}_L)^{-1} \quad (\text{B.8})$$

$$\hat{\mathbf{b}}^{[i]} = 2\hat{\lambda}_1^{[i-1]} \hat{\mathbf{C}}_b^{[i]} \mathbf{U}^T \text{Re}\{\hat{\mathbf{H}}^{[i] \text{H}} \tilde{\mathbf{r}}\} \quad (\text{B.9})$$

$$\hat{\lambda}_1^{[i]} = \frac{KNM}{\|\tilde{\mathbf{r}}\|^2 - 2\hat{\mathbf{b}}^{[i] \text{T}} \mathbf{U}^T \text{Re}\{\hat{\mathbf{H}}^{[i] \text{H}} \tilde{\mathbf{r}}\} + \hat{E}_b^{[i]} \hat{E}_h^{[i]}}. \quad (\text{B.10})$$

After initializing the messages as $\hat{\lambda}_1^{[0]} = KNM / \|\tilde{\mathbf{r}}\|^2$, $\hat{\mathbf{C}}_b^{[0]} = \mathbf{C}_{b_0}$, and $\hat{\mathbf{b}}^{[0]}$ as a realization drawn from the prior $p(\mathbf{b})$, equations (B.6) to (B.10) are iterated until the messages are converged. Furthermore, $\hat{\mathbf{C}}_b^{[i]}$ is calculated by adding a scaled identity matrix to the inverse prior $\mathbf{C}_{b_0}^{-1}$. Therefore, the eigenvectors \mathbf{U} do not

change throughout the iterations and can be precomputed based on the chosen prior \mathbf{C}_{b_0} .

To keep the notation concise, we refrain from explicitly writing iteration indices in the remainder of the paper, referring to the respective values after they are converged.

B.3.1 Detection

In order to detect the presence of a person, we need to distinguish between two nested models \mathcal{H}_0 and \mathcal{H}_1 , with corresponding likelihoods:

$$p(\tilde{\mathbf{r}}|\mathbf{b} = \mathbf{0}, \mathbf{h}_s = \mathbf{0}, \lambda, \mathcal{H}_0) = \mathcal{CN}(\tilde{\mathbf{r}}|\mathbf{0}, \lambda^{-1}\mathbf{I}_{KNM})$$

of an empty car, and

$$p(\tilde{\mathbf{r}}|\mathbf{b}, \mathbf{h}_s, \lambda, \mathcal{H}_1) = \mathcal{CN}(\tilde{\mathbf{r}}|\mathbf{H}\mathbf{U}\mathbf{b}, \lambda^{-1}\mathbf{I}_{KNM})$$

if a person is present. Since the ELBO, is a lower bound on the logarithmic model evidence $\mathcal{L}(q_j) \leq \ln p(\mathcal{H}_j|\tilde{\mathbf{r}})$, we approximate the log odds ratio as $\ln \frac{p(\mathcal{H}_1|\tilde{\mathbf{r}})}{p(\mathcal{H}_0|\tilde{\mathbf{r}})} \approx \mathcal{L}(q_1) - \mathcal{L}(q_0)$, where the ELBO is given as $\mathcal{L}(q_i) = \langle \ln p(\mathbf{b}, \mathbf{h}_s, \lambda, \mathcal{H}_i|\tilde{\mathbf{r}}) \rangle_{q_i(\mathbf{b}, \mathbf{h}_s, \lambda)} + \mathbb{H}(q_i)$ for $i \in \{0, 1\}$ [6, Ch. 10]. Since $q_0(\mathbf{b}, \mathbf{h}_s, \lambda) = q_0(\lambda)$ depends only on one parameter, we do not need an iterative update scheme and the ELBO $\mathcal{L}(q_0)$ is maximized by $q_0(\lambda) = \text{Ga}(\lambda | KNM, \|\tilde{\mathbf{r}}\|^2)$, resulting in $\hat{\lambda}_0 = KNM/\|\tilde{\mathbf{r}}\|^2$. Special considerations must be made regarding the improper prior $p(\lambda) \propto \lambda^{-1}$. Using a proper prior $p(\lambda) = \text{Ga}(\lambda | d, e)$ and taking the limit as $d, e \rightarrow 0$, the test decides for \mathcal{H}_1 if

$$(NM - 1) \ln \frac{\hat{\lambda}_1}{\hat{\lambda}_0} - \sum_{k=1}^K [\hat{\mathbf{h}}_{s,k}^H \mathbf{C}_{h_0,k}^{-1} \hat{\mathbf{h}}_{s,k} + \text{tr}(\mathbf{C}_{h_0,k}^{-1} \hat{\mathbf{C}}_{h,k})] - \frac{1}{2} [\hat{\mathbf{b}}^T \mathbf{C}_{b_0}^{-1} \hat{\mathbf{b}} + \text{tr}(\mathbf{C}_{b_0}^{-1} \hat{\mathbf{C}}_b)] + \mathbb{H}(q_1) - \mathbb{H}(q_0) > \gamma \quad (\text{B.11})$$

is larger than the detection threshold γ .

B.4 Results

To evaluate the performance of the devised algorithm, we consider the case of a single target sitting in a car. A raised-cosine pulse is transmitted every $T_{\text{rep}} = 0.1$ s during a measurement duration of 10 s using a bandwidth of $B = 500$ MHz and a roll-off factor of 0.5 at a centre frequency of $f_c = 6.5$ GHz corresponding to the UWB channel 5 in [22]. The forward and backward channels \mathbf{h}_f and \mathbf{h}_b are modeled with an LoS component with power E_{LoS} at a delay of $\tau_0 = 1$ m/c. The LoS component is followed by a diffuse multipath with exponentially decaying power delay profile with decay constant τ_f . Thus \mathbf{h}_f and \mathbf{h}_b are described by the covariance $C_{h_f}[n, n'] = C_{h_b}[n, n'] = [E_{\text{LoS}} + E_{\text{DM}} (1 + j2\pi\tau_f\Delta f(n - n'))^{-1}] e^{-j2\pi\tau_0\Delta f(n - n')}$. We choose $\tau_f = 20$ ns and $K_{\text{LoS}} = \frac{E_{\text{LoS}}}{E_{\text{DM}}} = 0.75$, since these values were observed by test measurements. If the delay of the LoS component τ_0 is known, the prior covariance \mathbf{C}_{h_0} can be calculated from $C_{h_f}[n, n']$. However, we do not assume the




Figure B.3: Detection performance of the VMP-based detector compared to an estimator-correlator (EC) and FFT detector on simulated data (top) and measurements (bottom).

distance to the target to be known a-priori. Therefore, we choose the covariance for the forward-backward channel to follow the shape of a gamma distribution $C_{fb}[n, n'] \propto (1 + j2\pi\tau_f\Delta f(n - n'))^{-a}$. We select $a = 1 + \frac{2\tau'_0}{\tau_f}$ such that the peak (in time domain) corresponds to the expected target distance $\tau'_0 = \frac{1m}{c}$. This choice of prior enhances the detection performance compared to a flat prior by incorporating the decay constant of the channel without requiring the exact target distance to be known. The effectiveness of this prior is confirmed by the results depicted in Fig. B.3. The typical respiratory frequency for an adult is between 9 and 21 breaths per minute [23], whereas the respiratory frequency of babies can go as high as 60 breaths per minute [12]. Therefore, we assume a prior covariance C_{b_t} with a rectangular double sided power spectral density $S_b(f) = \frac{1}{2(f_{b,max} - f_{b,min})}$ for $f_{b,min} \leq |f| \leq f_{b,max}$ and 0 elsewhere, and select $f_{b,min} = 9/60$ Hz and $f_{b,max} = 1$ Hz.

To evaluate the performance, a Monte-Carlo simulation with 10^5 runs was performed at each SNR $= \frac{\lambda \|b\|^2 \|h_s\|^2}{NM}$ from -30 dB to 10 dB in 0.5 dB steps considering a monostatic setup. Using numerical simulations, we set the threshold for all detectors such that a constant false alarm rate of $p_{FA} = 0.01$ was achieved at each SNR value. As comparison, we evaluated the detection performance of an estimator-correlator, which models \tilde{r} as a Gaussian process to incorporate MPCs into the detection [4], and an FFT-based detector which computes the FFT over the rows of the matrix $\mathbf{V}\mathbf{R}$ and compares the peak against a threshold. For the

FFT-based approach the channel is transformed to the time domain to concentrate the signal energy in the delay bin corresponding to the LoS component for easier detection. However, no MPCs are incorporated in the detection. The results are depicted in Fig. B.3. The proposed VMP-based algorithm has a better detection rate compared to the two other methods, even when using the modified prior which does not require a-priori knowledge of the distance to the target. Specifically, at an SNR of -20 dB, the VMP-based detection achieves a detection rate of approximately 0.95 while the estimator-correlator achieves 0.32 and the FFT detector achieves 0.05 in case the target distance is not known. This difference in performance can be explained by the different level of incorporation of the signal model into the detection: the VMP-based detector incorporates the signal model more rigidly in the detection compared to the estimator-correlator which only accounts for the correlation between different columns of the matrix $\tilde{\mathbf{R}}$, or the FFT-based detector which does not account for MPCs at all.

A measurement campaign including 34 participants (6 female and 28 male) was performed during which 177 minutes of data with a sample rate of $T_{\text{rep}} = 0.1$ s have been collected using a multistatic setup with one transmit and two receive antennas. The participants were instructed to sit motionless in the car while breathing normally. The measurements were performed using an M-sequence channel sounder in a Seat Leon and a Citroen Picasso. The measurement equipment including channel sounder, cables and connectors has been calibrated before the measurement and the same transmit pulse was used as in the simulations. The data was split in to non-overlapping chunks of 10 s length, for a total of 481 measurements. For each SNR value and sample, 100 independent noise realizations were added. The resulting detection performance is shown in Fig. B.3 using only data from one receive antenna ($K = 1$) and using data from both receive antennas ($K = 2$). The FFT-based approach performs significantly better on the measured data compared to the simulated data, since the respiratory motion of adults is closer to a periodic signal and has less randomness than the Gaussian process used in the simulations. Although the performance difference between the VMP-based approach and the other methods is less than in the simulations, the best detection probability is still achieved by the VMP-based detector.

B.5 Conclusion


We present a novel VMP-based approach to detect the presence of a person by their respiratory chest motion using UWB radar signals. The devised algorithm significantly outperforms the comparison methods on simulated as well as measured data in low-SNR conditions. The superior performance of the presented detection algorithm is achieved through the use of MPCs, which are shown to carry a significant amount of signal energy in small spaces, such as the interior of a car. Furthermore, the algorithm should be robust against deviations from a strictly periodic respiration pattern, such as pauses with different lengths between individual breaths since the respiratory motion $b_t(t)$ is modeled as a random process and not as a periodic function. Additionally, the devised algorithm can be used for contact-free vital sign estimation since the respiratory chest motion is estimated as part of the algorithm. However, the received signal $\tilde{\mathbf{r}}$ depends on the product of \mathbf{b} and \mathbf{h}_s which results in an ambiguity in the sign of the estimates $\hat{\mathbf{b}}$ and $\hat{\mathbf{h}}_s$.

Since many future cars will be equipped with UWB nodes, e.g. as part of the keyless entry system, the developed algorithm can provide occupancy sensing capabilities to these cars without the need and increased manufacturing costs of dedicated sensors.

Bibliography

- [1] Euro NCAP. (2021) Euro NCAP 2025 roadmap. Accessed: May 2021. [Online]. Available: <https://cdn.euroncap.com/media/30701/euroncap-roadmap-2025-v4-print.pdf>
- [2] Q. Xu, B. Wang, F. Zhang, D. S. Regani, F. Wang, and K. J. R. Liu, “Wireless AI in smart car: How smart a car can be?” *IEEE Access*, Mar. 2020, doi: 10.1109/ACCESS.2020.2978531.
- [3] A. S. Aghaei, B. Donmez, C. C. Liu, D. He, G. Liu, K. N. Plataniotis, H.-Y. W. Chen, and Z. Sojoudi, “Smart driver monitoring: When signal processing meets human factors: In the driver’s seat,” *IEEE Signal Process. Mag.*, Nov. 2016, doi: 10.1109/MSP.2016.2602379.
- [4] J. Möderl, F. Pernkopf, and K. Witrisal, “Car occupancy detection using UWB radar,” in *2021 18th Eur. Radar Conf.*, Apr. 5–7, 2022, doi: 10.23919/EuRAD50154.2022.9784505.
- [5] J. Winn, C. M. Bishop, and T. Jaakkola, “Variational message passing,” *Journal of Machine Learning Research*, vol. 6, no. 4, pp. 661–694, Apr. 2005.
- [6] C. M. Bishop, *Pattern recognition and machine learning*, 8th ed., M. Jordan, J. Kleinberg, and B. Schölkopf, Eds. New York, NY, USA: Springer-Verlag, 2009.
- [7] Z. Baird, I. Gunasekara, M. Bolic, and S. Rajan, “Principal component analysis-based occupancy detection with ultra wideband radar,” in *2017 IEEE 60th Int. Midwest Symp. on Circuits and Syst.*, Aug. 6–9, 2017, doi: 10.1109/MWSCAS.2017.8053237.
- [8] A. Ahmad, J. C. Roh, D. Wang, and A. Dubey, “Vital signs monitoring of multiple people using a FMCW millimeter-wave sensor,” in *2018 IEEE Radar Conf.*, Apr. 23–27, 2018, doi: 10.1109/RADAR.2018.8378778.
- [9] L. Anitori, A. de Jong, and F. Nennie, “FMCW radar for life-sign detection,” in *2009 IEEE Radar Conf.*, May 4–8, 2009, doi: 10.1109/RADAR.2009.4976934.
- [10] Y. Kilic, H. Wymeersch, A. Meijerink, M. J. Bentum, and W. G. Scanlon, “Device-free person detection and ranging in UWB networks,” *IEEE J. Sel. Topics Signal Process.*, Feb. 2014, doi: 10.1109/JSTSP.2013.2281780.
- [11] D. Yang, Z. Zhu, J. Zhang, and B. Liang, “The overview of human localization and vital sign signal measurement using handheld IR-UWB through-wall radar,” *Sensors*, May 2021, doi: 10.3390/s21020402.

- [12] U. Frey, M. Silverman, A. L. Barabási, and B. Suki, “Irregularities and power law distributions in the breathing pattern in preterm and term infants,” *J. Appl. Physiol.*, Sep. 1998, doi: 10.1152/jappl.1998.85.3.789.
- [13] Y. Ma, Y. Zeng, and V. Jain, “CarOSense: Car occupancy sensing with the ultra-wideband keyless infrastructure,” *Proc. ACM Interact. Mobile Wearable Ubiquitous Technol.*, Sep. 2020, doi: 10.1145/3411820.
- [14] M. Alizadeh, H. Abedi, and G. Shaker, “Low-cost low-power in-vehicle occupant detection with mm-wave FMCW radar,” in *IEEE SENSORS*, Oct. 27–30, 2019, doi: 10.1109/SENSORS43011.2019.8956880.
- [15] G. D. Durgin, *Space-Time Wireless Channels*, 1st ed., T. S. Rappaport, Ed. Upper Saddle River, NJ, USA: Prentice-Hall, 2002.
- [16] D. Arnitz, U. Muehlmann, and K. Witrisal, “Wideband characterization of backscatter channels: Derivations and theoretical background,” *IEEE Trans. Antennas Propag.*, Jan. 2012, doi: 10.1109/TAP.2011.2167923.
- [17] G. E. Kinkelund, C. N. Manchon, L. P. B. Christensen, E. Riegler, and B. H. Fleury, “Variational message-passing for joint channel estimation and decoding in MIMO-OFDM,” in *2010 IEEE Global Telecommun. Conf.*, Dec. 6–10, 2010, doi: 10.1109/GLOCOM.2010.5683839.
- [18] E. Riegler, G. E. Kinkelund, C. N. Manchon, M. A. Badiu, and B. H. Fleury, “Merging belief propagation and the mean field approximation: A free energy approach,” *IEEE Trans. Inf. Theory*, Jan. 2013, doi: 10.1109/tit.2012.2218573.
- [19] C. Zhang, J. Bütepage, H. Kjellström, and S. Mandt, “Advances in variational inference,” *IEEE Trans. Pattern Anal. Mach. Intell.*, Aug. 2019, doi: 10.1109/TPAMI.2018.2889774.
- [20] T. Minka, “Divergence measures and message passing,” Microsoft Res., Tech. Rep., 2005, accessed: Sep. 2022. [Online]. Available: <https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.361.9105&rep=rep1&type=pdf>
- [21] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, “Variational inference: A review for statisticians,” *J. Amer. Statistical Assoc.*, Jul. 2017, doi: 10.1080/01621459.2017.1285773.
- [22] *IEEE Standard for Low-Rate Wireless Networks*, IEEE Std. 802.15.4-2020, Jul. 2020, (Revision of IEEE Std 802.15.4-2015).
- [23] S. H. Fairclough and L. J. M. Mulder, “Psychophysiological processes of mental effort investment,” in *How motivation affects cardiovascular response: Mechanisms and applications*, R. A. Wright and G. H. E. Gendolla, Eds. Amer. Psychological Assoc., 2012, pp. 61–76.



UWBCarGraz Dataset for Car Occupancy Detection using Ultra-Wideband Radar

Jakob Möderl, Stefan Posch, Franz Pernkopf, and Klaus Witrisal

*2024 IEEE Radar Conference, Denver, CO, USA, May 04–10, 2024,
doi: 10.1109/RadarConf2458775.2024.10549280*

Abstract—We present a data-driven car occupancy detection algorithm using ultra-wideband radar based on the ResNet architecture. The algorithm is trained on a dataset of channel impulse responses obtained from measurements at three different activity levels of the occupants (i.e. breathing, talking, moving). We compare the presented algorithm against a state-of-the-art car occupancy detection algorithm based on variational message passing (VMP). Our presented ResNet architecture is able to outperform the VMP algorithm in terms of the area under receiver operating curve (AUC) at low signal-to-noise ratios (SNRs) for all three activity levels of the target. Specifically, for an SNR of -20 dB our ResNet architecture achieves an AUC of 0.91 while the VMP detector only achieves an AUC of 0.87 if the target is sitting still and breathing naturally. The difference in performance for the other activities is similar. Furthermore, to facilitate the implementation in the onboard computer of a car, we train a collection of different ResNet architectures to find a balance between the detection performance and computational complexity. The UWBCarGraz dataset used to train and evaluate the algorithm is openly accessible.

This research was partly funded by the Austrian Research Promotion Agency (FFG) within the project SEAMAL Front (project number: 880598). Furthermore, the financial support by the Christian Doppler Research Association, the Austrian Federal Ministry for Digital and Economic Affairs and the National Foundation for Research, Technology and Development is gratefully acknowledged.

C.1 Introduction

Radar will become a pervasive technology through the advent of civilian applications e.g. in assisted living for vital-sign monitoring [1, 2] and human activity recognition [3]. Similarly, radar can be deployed in the passenger cabin of a car, e.g. to detect occupants to prevent the confinement and endangerment of small children [4–7]. New generation car models are equipped with ultra-wideband (UWB) nodes for keyless car access, that can be used for opportunistic sensing of occupants [7–9]. The signal received by the radar contains the radar response from the target, but also includes additive noise and (strong) clutter [8, 9]. To distinguish the target signal from clutter, most works focus on detecting variations of the channel impulse response (CIR) over time, which correspond to movements of the target [8–13]. This movement can range from a subtle chest movement due to the respiration of the target to rather larger movements, such as limb movements or a repositioning of the torso.

Several model-driven [8, 9, 12–16] and data-driven [7, 17–20] approaches to detect car occupants can be found in the literature which utilize either UWB or frequency-modulated continuous-wave (FMCW) radar at various carrier frequencies. However, none of these works have made the dataset used to train and/or evaluate their algorithms publicly available. This makes it difficult to compare and assess the results obtained by different methods in a fair and unbiased manner. This is exacerbated for data-driven methods, which cannot be properly reproduced and validated without the data used to train the network.

In previous works, we focused on modeling and detecting the case of small movement of the target due to respiration [8, 9]. For such small movements, the CIRs received from the target over time can be (approximately) factorized into a time-independent channel multiplied with the time-varying movement amplitude of the target. However, this signal model is only valid for small body movements, but not for larger movements such as limb or torso movements. Furthermore, the prior chosen to model the respiratory motion is rather general and might not represent the true physiological motion accurately, e.g. when the target is speaking. In this paper, we address these issues using a data-driven approach, where the statistical relations of the detection problem are learned by a ResNet architecture [21]. Hence, the two main contributions of this paper are

- We implement, train, and evaluate a collection of ResNet architectures for car occupancy detection to achieve a balance between detection performance and computational complexity. Our ResNet classifier outperform state-of-the-art model-based methods, such as the variational message passing (VMP) detector [9], for car occupancy detection in terms of the area under receiver operating curve (AUC) at low signal-to-noise ratios (SNRs).
- We present the UWBCarGraz dataset [22], an open dataset for car occupancy detection using UWB radar.

C.2 System Model

We consider a radar system which repeatedly transmits a pulse $s(\tau)$ with center frequency f_c and bandwidth B at times $t = mT_{\text{st}}$, $m \in \{0, 1, \dots, M - 1\}$ with

repetition interval T_{st} , where τ denotes “fast time” in the order of nanoseconds and t denotes “slow time” in the order of seconds. The signal propagates through the environment where it is reflected by the target as well as all other objects in the environment. At each time instance $t = mT_{\text{st}}$, N samples with sampling interval T_{ft} are obtained. Assuming that $T_{\text{st}} \gg NT_{\text{ft}}$ is chosen large enough such that successive repetitions of the transmit signal do not interfere with each other, the received signal $r_m[n]$, $n \in \{0, 1, \dots, N-1\}$ at repetition m is modeled as

$$r_m[n] = (h(\cdot, mT_{\text{st}}) * s(\cdot))(nT_{\text{ft}}) + v(nT_{\text{ft}}, mT_{\text{st}}) \quad (\text{C.1})$$

where $(h(\cdot) * s(\cdot))(nT_{\text{ft}})$ denotes the (fast time) convolution of the functions h and s evaluated at nT_{ft} , $h(\tau, t)$ is the time-varying CIR and $v(\tau, t)$ is an additive noise process which is assumed to be white and independent across different repetitions m . We use a multipath propagation model [23, 24] to describe the CIR $h(\tau, t)$ as a sum of L multipath components. Of these L multipath components, K components interact with the target and $L - K$ components interact only with the environment. Assuming a static environment, we express the CIR $h(\tau, t)$ as

$$h(\tau, t) = \sum_{k=1}^K \alpha_k(t) \delta(\tau - \tau_k(t)) + \sum_{k=K+1}^L \alpha_k \delta(\tau - \tau_k) \quad (\text{C.2})$$

where $\delta(\cdot)$ denotes the Dirac-delta function, α_k and τ_k denote the amplitude and delay of the k -th multipath component, which are time-varying for $k \leq K$, i.e. for all multipath components associated with the target but not for the multipath components associated with the static environment. Inserting (C.2) into (C.1) we get

$$\begin{aligned} r_m[n] = & \sum_{k=1}^K \alpha_k(mT_{\text{st}}) s(nT_{\text{ft}} - \tau_k(mT_{\text{st}})) \\ & + \sum_{k=K+1}^L \alpha_k s(nT_{\text{ft}} - \tau_k) + v(nT_{\text{ft}}, mT_{\text{st}}). \end{aligned} \quad (\text{C.3})$$

Let $\mathbf{r}_m = [r_m[0] \ r_m[1] \ \dots \ r_m[N-1]]^T$ denote the vector of the received signal samples at time m . We stack the signals received at different times m into a matrix $\mathbf{R} \in \mathbb{C}^{N \times M}$

$$\mathbf{R} = [\mathbf{r}_0 \ \mathbf{r}_1 \ \dots \ \mathbf{r}_{M-1}]. \quad (\text{C.4})$$

From (C.3) it follows that the columns \mathbf{r}_m of \mathbf{R} can be decomposed into a constant part $\bar{\mathbf{r}} = \frac{1}{M} \sum_{m=1}^M \mathbf{r}_m$ and a time-varying part $\tilde{\mathbf{r}}_m = \mathbf{r}_m - \bar{\mathbf{r}}$ which includes the variations in the CIR over time due to the target plus noise. If no target is present, the columns $\tilde{\mathbf{r}}_m$ will contain only noise. Hence, we focus on detecting if the mean-removed received signal

$$\tilde{\mathbf{R}} = [\tilde{\mathbf{r}}_0 \ \tilde{\mathbf{r}}_1 \ \dots \ \tilde{\mathbf{r}}_{M-1}] \quad (\text{C.5})$$

contains only white noise or, if variations due to a target are present.¹

¹Note, that also the static part $\bar{\mathbf{r}}$ potentially depends on the target. However, to avoid false positives due to e.g. a suitcase put in the car we focus solely on the time-varying part $\tilde{\mathbf{R}}$ in the detection process.




Figure C.1: Illustration of the measurement setup. The participant positions are indicated with gray circles.

C.3 UWBCarGraz Dataset Description

A measurement campaign was conducted at Graz University of Technology with 34 adult participants (6 female, 28 male). The participants are volunteers recruited from the staff and students of the university. To gather CIRs, a maximum length-sequence (MLS) correlative channel sounder [25] with a nominal frequency range of 3.8–10.2 GHz was used to measure the CIR between one transmit and two receive antennas with a repetition interval of $T_{\text{st}} = 0.1$ s. Two pairs of antennas were placed under the ceiling of the passenger cabin near the A- and B-pillars. One antenna of each pair was receiving while the transmit signal was connected to an RF-switch such that either of the remaining two antennas could be used as transmit antenna. The measurement setup is illustrated in Figure C.1. The equipment, including channel sounder, cables and connectors, was calibrated before the measurement except for the antennas. Dipole antennas with circular patches have been used [26, Figure B.5b]. The pairs of antennas were placed such that the zeros in the radiation patterns of the antennas pointed towards each other in order to minimize the mutual coupling between them. In the post processing, the data was multiplied in the frequency domain with a raised cosine pulse with center frequency $f_c = 6.5$ GHz, bandwidth $B = 500$ MHz and roll-off factor of $\beta = 0.5$, corresponding to the UWB Channel 5 in [27]. A second variant of the dataset with a bandwidth of 5 GHz centered around $f_c = 7$ GHz is also available. Additionally, the expansion of the chest in each breathing cycle of the participants was recorded with an elastic belt [28]. Therefore, the UWBCarGraz dataset [22] can be used for breathing motion estimation as well. Further details about the calibration and data pre-processing can be found in the technical document accompanying the dataset.

The participants were seated in the front or rear seat on the passenger side or in the rear middle seat as shown in Figure C.1. The participants were not seated on the driver side for practical reasons. However, due to the symmetry of the setup we do not expect the CIRs recorded with people sitting in the driver side to differ substantially from the ones recorded with people on the passenger seats. In order to record different types of movements, the participants were asked to perform three different activities for approximately 2 minutes each. The data from each activity was then segmented into non-overlapping samples with 10 s length. The activities are

Table C.1: Bio-metrics of the participants.

	min	max	median	unit
height	158	190	178.5	cm
weight	45	158	72	kg
age	22	64	29	years

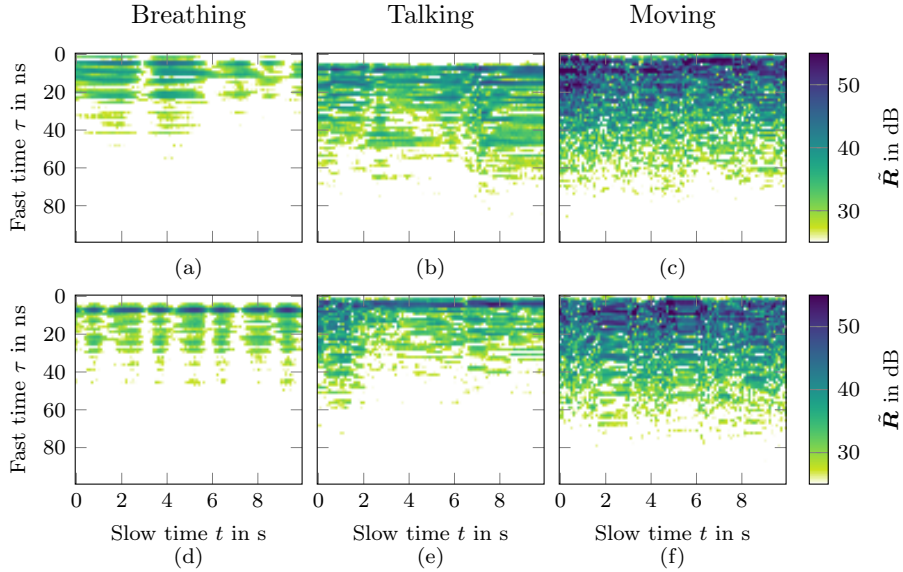


Figure C.2: Example signals from the dataset. The plots show the magnitude of the received signal \tilde{R} in dB for the breathing activity (a, d), talking activity (b, e) and moving activity (c, f). The x-axis is the measurement time in seconds (“slow time”) whereas the y-axis is the propagation delay in nanoseconds (“fast time”). The color scale is clipped so that all values ≥ 55 dB map to dark blue.

- (i) breathing: sitting still without intentional movements while breathing naturally,
- (ii) talking: speaking out loud either a freely improvised monologue or by reading from a book, and
- (iii) moving: looking around and searching for small objects hidden inside the car.

Additionally, 5 minutes of data of the empty car was also recorded. Measurements have been performed with two different cars: a hatchback compact car (Car 1, Seat Leon) and a minivan (Car 2, Citroën C4 Picasso). Out of the 34 participants, 9 were recorded twice, once in each car, while the remaining 25 participants were recorded only in one of the cars for a total of 21 participants in the hatchback compact car and 22 participants in the minivan. Table C.1 list the min, max and median of the participants height, weight and age. It is evident from Table C.1, that a diverse group participated in the measurement campaign.

Examples of the time-varying part of the received signal $\tilde{\mathbf{R}}$ of the three different activity levels are shown in Figure C.2. Each column of $\tilde{\mathbf{R}}$ corresponds to the time-varying part of the CIR at time $t = nT_{\text{st}}$ and each row corresponds to the evolution of a tap of the CIR over slow time t . In the breathing case (i), the columns of $\tilde{\mathbf{R}}$ have approximately the shape of a line-of-sight component followed by a dense multipath profile. The magnitude of each column changes according to the periodic chest movement of the target. This is in agreement with the model of [8,9]. In the talking case (ii), the chest motion is not periodic anymore and variations can be observed across different propagation delays. Due to larger chest movements, the amplitude of the variations in the CIR (≈ 55 dB maximum) is larger compared to the breathing case (≈ 50 dB maximum). In the movement case (iii), no clear structure is visible in the data and the maximum amplitude of the variations (≈ 65 dB) is once again larger compared to the talking case due to larger movements. The colorscale in Figure C.2 is clipped such that all values ≥ 55 dB map to dark blue.

C.4 Data-driven Occupancy Detection

C.4.1 ResNet Architecture

We used two different neural-network architectures based on the ResNet architecture [21]. For each of these architectures, several variants were evaluated with varying computational complexity.

For the first architecture, we used a two-dimensional convolutional ResNet. Each sample $\tilde{\mathbf{R}} \in \mathbb{C}^{N \times M}$ is transformed to a sample $\mathbf{R}_{\text{input}} \in \mathbb{R}^{N \times M \times 2}$ by adding a new dimension for the real and imaginary part of the complex-valued sample $\tilde{\mathbf{R}}$. The network architecture consists of a repeated structure of pairs of 2D-convolutional layers and batch normalization layers, with skip connections running in parallel to every pair of convolutional layers and batch normalization layers after the initial convolutions. We denote such a block as “ResNet block”. After N_{double} ResNet blocks, the number of convolutional filters (i.e. channels) doubles. A total of N_{total} ResNet blocks is used. To preserve the size, we used *same padding*. The network ends with an average pooling layer across each channel, followed by a flattening and a dense output layer. This general architecture is shown in Figure C.3a.

The berathing activity results in the signal with the lowest energy [29]. Hence, it is the hardest to detect. For this activity, the received signal is mostly characterized by periodic movement along slow time [8,9]. To detect these periodic movements with lower computational effort than the 2D-convolutions used in the first architecture, we also investigate one-dimensional convolutional ResNets. The rows of $\tilde{\mathbf{R}} \in \mathbb{C}^{N \times M}$ corresponding to the samples over slow time t at one fast time instance τ are considered as independent channels and $\tilde{\mathbf{R}}$ is transformed to a sample $\mathbf{R}_{\text{input}} \in \mathbb{R}^{2N \times M}$ by stacking the real and imaginary parts on top of each other, resulting in a 1D slow time signal of with $2N$ channels. The network follows a similar structure as the two-dimensional ResNet, with a repeated structure of 1D-convolutional layers along slow time and batch normalization layers, with skip connections running in parallel to every pair of convolutional layers. Like the 2D architecture, the network ends with an average pooling layer across each channel followed by a flattening and a dense output layer. Figure C.3b illustrates this architecture.

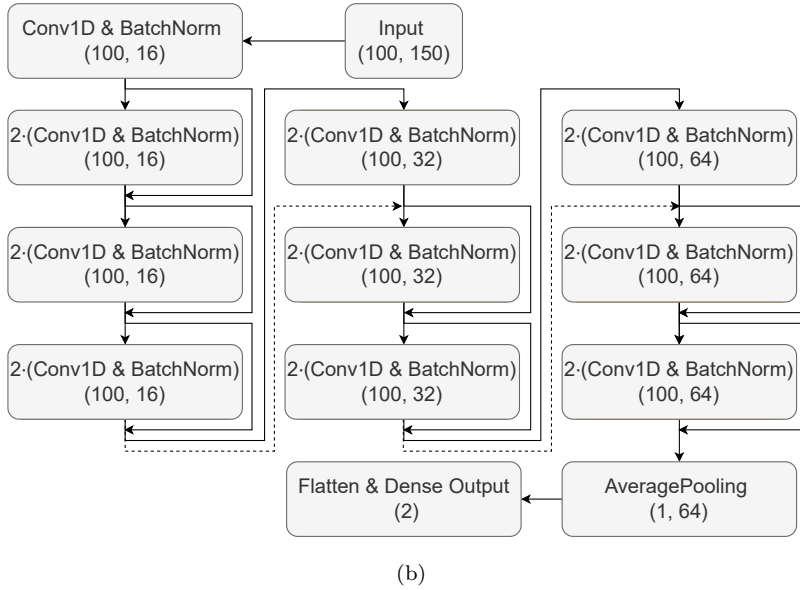
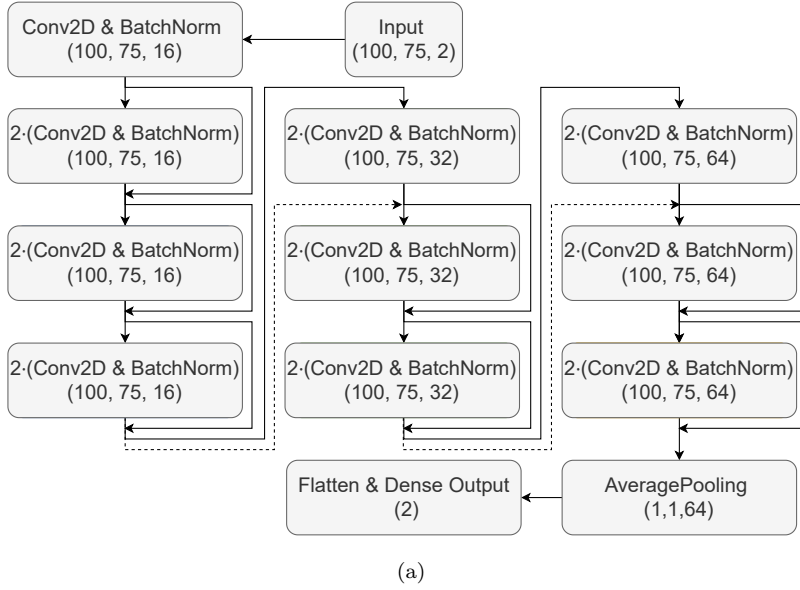


Figure C.3: Visualization of the ResNet-architectures for the (a) 2D-A and (b) 1D-B ResNet variants. For details of the architectures see Table C.2. Connections and skip connections are denoted with arrows and projection skip connections are denoted with dotted arrows. The numbers in brackets denote the output dimensions of a block.

Table C.2: Evaluated model architectures.

Variant	Initial filters	N_{double}	N_{total}	# Parameters	FLOPs
1D-A	32	3	12	$1.5 \cdot 10^6$	$3.0 \cdot 10^8$
1D-B	16	3	9	$1.0 \cdot 10^5$	$2.0 \cdot 10^7$
1D-C	16	2	6	$6.8 \cdot 10^4$	$1.3 \cdot 10^7$
1D-D	16	1	3	$3.5 \cdot 10^4$	$6,9 \cdot 10^6$
1D-E	8	1	3	$1.0 \cdot 10^4$	$2.1 \cdot 10^6$
2D-A	16	3	9	$2.7 \cdot 10^5$	$4.0 \cdot 10^9$
2D-B	16	2	6	$1.7 \cdot 10^5$	$2.6 \cdot 10^9$
2D-C	8	2	6	$4.4 \cdot 10^4$	$6.5 \cdot 10^8$
2D-D	8	2	4	$1.1 \cdot 10^4$	$1.6 \cdot 10^8$
2D-E	4	2	4	$2.9 \cdot 10^3$	$4.1 \cdot 10^7$

Table C.3: Number of Samples in the UWBCarGraz dataset.

	Car 1 (Seat Leon)			Car 2 (Citroën C4 Picasso)		
	Train	Val.	Test	Train	Val.	Test
Breathing	368	144	0	0	409	150
Talking	367	145	0	0	406	150
Moving	380	161	0	0	410	150
Empty	0	0	0	66	100	20

Both of these models were used with different variations in depth (number of layers) and width (number of filter maps), leading to different model complexities. Table C.2 lists the parameters for each model variant along with the number of trainable parameters and the floating point operations (FLOPs) of a single forward pass.

C.4.2 Data Augmentation and Training

The data were split in training data and validation/test data between the different cars. For the training data, only the measurements from Car 1 are used. The data from Car 2 are used as validation and test data, with the last 150 samples of each class comprising the test set and the rest of the data from Car 2 being used for the validation set. While the validation set was used to govern early stopping during training, the test set was only used after the architectures have been designed and the networks have been trained. With this split, one of the two cars and 13 participants which were unseen during training are used only in the validation set and the test set. Table C.3 lists the exact number of samples of each car used for training, validation and test. For the empty car, samples from Car 2 are available only.

The UWBCarGraz dataset contains fewer samples of the empty car compared to the occupied case (see Table C.3). To account for this imbalance and to achieve a balanced dataset for training and testing, the samples of the occupied class are used 200 times each whereas the samples of the empty class are used 3000 times each. The resulting augmented dataset is mostly balanced, e.g. the test set contains $450 \times 200 = 90\,000$ samples of the occupied class and $20 \times 3000 = 60\,000$ samples of

the empty class. Overfitting due to the reuse of samples, especially in the empty case, is addressed by data augmentation. After the data augmentation described in the following paragraph, the samples of the empty class are virtually indistinguishable from white noise due to the low energy of these samples.

Due to the high dynamic range and low noise floor of the laboratory equipment used in gathering the data, the detection process is trivial without data augmentation. However, in a real scenario the SNR of the UWB radar will be lower compared to our laboratory equipment. Furthermore, only adults were measured in the dataset which have a large radar cross section and chest movement compared to infants. Thus, the expected variation in the CIR is much lower for an infant which further decreases the SNR. To account for the lower SNR expected in a practical scenario, while also accounting for the different signal amplitudes at different activity levels, we define the $\text{SNR} = E_s / \|\mathbf{V}\|_2^2$ with respect to the median signal energy of the breathing case $E_s = \text{median}_{\tilde{\mathbf{R}} \in \mathcal{R}_b} (\|\tilde{\mathbf{R}}\|_2^2)$, where \mathcal{R}_b denote the set of all samples of the breathing activity, and the matrix operator $\|\cdot\|_2$ denotes the Frobenius norm, such that $\|\tilde{\mathbf{R}}\|_2^2$ equals the energy of the received signal, and $\mathbf{V} \in \mathbb{C}^{N \times M}$ denotes a matrix of samples of the additive noise process $v(nT_{\text{ft}}, mT_{\text{st}})$, sampled at $n \in \{0, 1, \dots, N-1\}$ and $m \in \{0, 1, \dots, M-1\}$. Each copy of each sample is augmented by adding Gaussian noise such that different SNRs drawn uniformly from the interval $[-30 \text{ dB}, 0 \text{ dB}]$ are achieved. Specifically, circularly-symmetric complex white Gaussian noise \mathbf{V} with real and imaginary sample variance $\sigma^2 = E_s / (2MN \cdot 10^{\frac{\text{SNR}}{10}})$ is added to each sample instance $\tilde{\mathbf{R}}$, such that in case this level of noise is added to each sample of the breathing activity, the desired SNR is achieved as the median of these augmented breathing samples. During the evaluation, we evaluate the full test set once for each $\text{SNR} \in \{-10 \text{ dB}, -11 \text{ dB}, \dots, -40 \text{ dB}\}$. Additionally, each augmented sample is normalized to unit energy before being fed into the ResNet.

C.5 Results

C.5.1 Classification Result

While the task of the ResNet is a binary decision between the car being empty or occupied, we evaluate the 3 different activities breathing, talking and moving separately due to the differences in scale and structure of the data. We measure the detection performance using the area under receiver operating curve (AUC), which is a threshold independent measure for a binary classifier. Figure C.4 shows the AUC on the test data as a function of the SNR for the three different activities. The results are compared against a state-of-the-art model-driven approach based on VMP [9], which was shown to outperform other model-driven approaches such as the estimate-correlator approach of [8] or an fast Fourier transform (FFT)-based detector [9]. As evident from Figure C.4, our ResNet is able to achieve a higher AUC than the VMP detector in all three scenarios and over all tested SNRs. The performance difference is the highest for the movement scenario and somewhat smaller for the breathing and talking scenario. The assumptions of the model-driven approach are (approximately) fulfilled in the breathing scenario, but not in the talking or moving scenarios. Nevertheless, even in the breathing scenario our ResNet-based approach outperforms the VMP approach. We attribute this to model mismatch and the difficulty of specifying the correct prior for the physiological movement of the chest




Figure C.4: AUC over SNR for the 2D-A ResNet variant for the three activities compared to the VMP-detector [9].

due to breathing in the VMP approach. Because this biological process is hard to model, the VMP approach assumes that the chest motion is a Gaussian process with a specified power spectral density. However, this prior is rather general. As indicated by the better detection performance, the ResNet seems to be able to (implicitly) estimate either a better fitting model or a more informative prior (or both) from the training data. The aforementioned model mismatch of the VMP detector is higher for the stronger and more erratic movements of the talking and moving scenarios, which explains the higher performance difference between the ResNet and VMP detectors in these cases.

C.5.2 Architecture Optimization

Since the detection algorithm is supposed to run on the resource-constraint onboard computer of a car, the computational efficiency of the ResNet forward pass is of interest. We performed an ablation study by gradually decreasing the number of channels, layers and filter banks in each layer of the ResNet for both the 1D and 2D variants. Figure C.5 shows the resulting AUC values of the ResNets as a function of the number of FLOPs required for the forward pass. Since the three different scenarios have different signal energies on average, we evaluate the AUC at SNR = -20 dB for the breathing subclass, at SNR = -24 dB for the talking subclass and at SNR = -30 dB for the moving subclass. The 2D-ResNet variants generally require more FLOPs than the 1D-ResNet variants but achieve a slightly better performance. The best tradeoff between the computational complexity is achieved by the 1D-D ResNet variant, which still outperforms the VMP-based approach but requires less than 10^7 FLOPs. Note, that the VMP approach is an iterative approach. Thus, the number of required FLOPs is not constant and depends on the number of iterations needed to converge.




Figure C.5: Detection performance of the different ResNet-architectures as function of the computational complexity of the forward-pass. Results are taken at SNR = -20 dB for the breathing activity (black), SNR = -24 dB for the talking activity (red) and SNR = -30 dB for the moving activity (cyan). Best viewed in color.

C.6 Conclusion

We present a ResNet-based car occupancy detection algorithm. The presented data-driven algorithm is shown to outperform state-of-the art model-driven occupancy detection algorithms, such as the VMP-based approach of [9], in terms of AUC at low SNRs for all three activities breathing, talking and moving. In order to facilitate the implementation on the onboard computer of a car we evaluated several ResNet architectures to reduce the number of FLOPs needed for a single forward pass of the network. The resulting 1D-ResNet architecture requires less than 10^7 FLOPs for a forward pass while still outperforming the VMP-based approach.

The UWBCarGraz dataset [22] used to train and evaluate our detection algorithm is openly accessible to facilitate an easy and fair comparison between different approaches. This is in contrast to existing literature on UWB-radar based car occupancy detection for which no open dataset exists to the author’s knowledge.

Bibliography

- [1] G. Paterniani, D. Sgreccia, A. Davoli, G. Guerzoni, P. Di Viesti, A. C. Valenti, M. Vitolo, G. M. Vitetta, and G. Boriani, “Radar-based monitoring of vital signs: A tutorial overview,” *Proc. IEEE*, Mar. 2023, doi: 10.1109/JPROC.2023.3244362.
- [2] Z. Xie, B. Zhou, X. Cheng, E. Schoenfeld, and F. Ye, “VitalHub: Robust, non-touch multi-user vital signs monitoring using depth camera-aided UWB,” in *2021 IEEE 9th Int. Conf. Healthcare Inform.*, Aug. 09–12, 2021, doi: 10.1109/ICHI52183.2021.00056.

- [3] S. Yang, J. L. Kernec, O. Romain, F. Fioranelli, P. Cadart, J. Fix, C. Ren, G. Manfredi, T. Letertre, I. D. H. Sáenz, J. Zhang, H. Liang, X. Wang, G. Li, Z. Chen, K. Liu, X. Chen, J. Li, X. Wu, Y. Chen, and T. Jin, "The human activity radar challenge: Benchmarking based on the 'Radar Signatures of Human Activities' dataset from Glasgow University," *IEEE J. Biomed. Health Inform.*, Apr. 2023, doi: 10.1109/JBHI.2023.3240895.
- [4] Euro NCAP, *Euro NCAP 2025 Roadmap*, 2021, accessed: May 2021. [Online]. Available: <https://cdn.euroncap.com/media/30701/euroncap-roadmap-2025-v4-print.pdf>
- [5] Q. Xu, B. Wang, F. Zhang, D. S. Regani, F. Wang, and K. J. R. Liu, "Wireless AI in smart car: How smart a car can be?" *IEEE Access*, Mar. 2020, doi: 10.1109/ACCESS.2020.2978531.
- [6] A. S. Aghaei, B. Donmez, C. C. Liu, D. He, G. Liu, K. N. Plataniotis, H.-Y. W. Chen, and Z. Sojoudi, "Smart driver monitoring: When signal processing meets human factors: In the driver's seat," *IEEE Signal Process. Mag.*, Nov. 2016, doi: 10.1109/MSP.2016.2602379.
- [7] Y. Ma, Y. Zeng, and V. Jain, "CarOSense: Car occupancy sensing with the ultra-wideband keyless infrastructure," *Proc. ACM Interact. Mobile Wearable Ubiquitous Technol.*, Sep. 2020, doi: 10.1145/3411820.
- [8] J. Möderl, F. Pernkopf, and K. Witrisal, "Car occupancy detection using UWB radar," in *2021 18th Eur. Radar Conf.*, Apr. 5–7, 2022, doi: 10.23919/EuRAD50154.2022.9784505.
- [9] J. Möderl, E. Leitinger, F. Pernkopf, and K. Witrisal, "Variational message passing-based respiratory motion estimation and detection using radar signals," in *2023 IEEE Int. Conf. on Acoust., Speech and Signal Process.*, Jun. 4–10, 2023, doi: 10.1109/ICASSP49357.2023.10095725.
- [10] B. Schleicher, I. Nasr, A. Trasser, and H. Schumacher, "IR-UWB radar demonstrator for ultra-fine movement detection and vital-sign monitoring," *IEEE Trans. Microw. Theory Techn.*, May 2013, doi: 10.1109/TMTT.2013.2252185.
- [11] M. Leib, W. Menzel, B. Schleicher, and H. Schumacher, "Vital signs monitoring with a UWB radar based on a correlation receiver," in *Proc. 4th Eur. Conf. Antennas and Propag.*, Barcelona, Spain, Apr. 12–16, 2010, pp. 1–5.
- [12] A. Lazaro, M. Lazaro, R. Villarino, and D. Girbau, "Seat-occupancy detection system and breathing rate monitoring based on a low-cost mm-Wave radar at 60 GHz," *IEEE Access*, Aug. 2021, doi: 10.1109/ACCESS.2021.3105390.
- [13] N. Munte, A. Lazaro, R. Villarino, and D. Girbau, "Vehicle occupancy detector based on FMCW mm-Wave radar at 77 GHz," *IEEE Sensors J.*, Dec. 2022, doi: 10.1109/JSEN.2022.3218454.
- [14] Z. Baird, I. Gunasekara, M. Bolic, and S. Rajan, "Principal component analysis-based occupancy detection with ultra wideband radar," in *2017 IEEE 60th Int. Midwest Symp. Circuits and Syst.*, Aug. 6–9, 2017, doi: 10.1109/MWS-CAS.2017.8053237.

- [15] H. Song and H.-C. Shin, “Single-channel FMCW-radar-based multi-passenger occupancy detection inside vehicle,” *Entropy*, Nov. 2021, doi: 10.3390/e23111472.
- [16] A. R. Diewald, J. Landwehr, D. Tatarinov, P. Di Mario Cola, C. Watgen, C. Mica, M. Lu-Dac, P. Larsen, O. Gomez, and T. Goniva, “RF-based child occupation detection in the vehicle interior,” in *17th Int. Radar Symp.*, May 10–12, 2016, doi: 10.1109/IRS.2016.7497352.
- [17] M. Alizadeh, H. Abedi, and G. Shaker, “Low-cost low-power in-vehicle occupant detection with mm-wave FMCW radar,” in *IEEE Sensors J.*, Oct. 27–30, 2019, doi: 10.1109/SENSOR43011.2019.8956880.
- [18] A. K. Sriranga, Q. Lu, and S. Birrell, “Novel radar based in-vehicle occupant detection using convolutional neural networks,” in *IEEE Symp. Series Comput. Intell.*, Dec. 4–7, 2022, doi: 10.1109/SSCI51031.2022.10022146.
- [19] P. E. Numan, H. Park, J. Lee, and S. Kim, “Machine learning-based joint vital signs and occupancy detection with IR-UWB sensor,” *IEEE Sensors J.*, Apr. 2023, doi: 10.1109/JSEN.2023.3247728.
- [20] S.-Y. Kwon and S. Lee, “In-vehicle seat occupancy detection using ultra-wideband radar sensors,” in *23rd Int. Radar Symp.*, Sep. 12–14, 2022, doi: 10.23919/IRS54158.2022.9905064.
- [21] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *2016 IEEE Conf. on Comput. Vision and Pattern Recognit.*, Jun. 27–30, 2016, doi: 10.1109/CVPR.2016.90.
- [22] J. Möderl, S. Posch, F. Pernkopf, and K. Witrisal, “UWB CarGraz dataset,” Graz University of Technology, Nov. 2023, doi: 10.3217/2gx2m-pt043.
- [23] N. Michelusi, U. Mitra, A. F. Molisch, and M. Zorzi, “UWB sparse/diffuse channels, Part I: Channel models and bayesian estimators,” *IEEE Trans. Signal Process.*, Jun. 2012, doi: 10.1109/TSP.2012.2205681.
- [24] A. F. Molisch, *Wireless Communications*, 2nd ed. Hoboken, NJ, USA: Wiley, 2011.
- [25] J. Sachs, R. Herrmann, M. Kmec, M. Helbig, and K. Schilling, “Recent advances and applications of M-sequence based ultra-wideband sensors,” in *2007 IEEE Int. Conf. Ultra-Wideband*, Sep. 24–26, 2007, doi: 10.1109/ICUWB.2007.4380914.
- [26] C. Krall, “Signal processing for ultra wideband transceivers,” PhD Thesis, Signal Process. and Speech Commun. Lab., Graz Univ. of Technol., Graz, Austria, 2008.
- [27] *IEEE Standard for Low-Rate Wireless Networks*, IEEE Std. 802.15.4-2020, Jul. 2020, (Revision of IEEE Std 802.15.4-2015).
- [28] Vernier Science Education, *Datasheet for Go Direct Respiration Belt*, accessed: Sep. 2023. [Online]. Available: <https://www.vernier.com/files/manuals/gdx-rb/gdx-rb.pdf>

- [29] Z. Xie, B. Zhou, and F. Ye, “Signal quality detection towards practical non-touch vital sign monitoring,” in *Proc. 12th ACM Conf. Bioinf., Comput. Biol., Health Inform.*, Aug. 1–4, 2021, doi: 10.1145/3459930.3469526.



Fast Variational Block-Sparse Bayesian Learning

Jakob Möderl, Franz Pernkopf, Klaus Witrisal, and Erik Leitinger

*Submitted to IEEE Transaction on Signal Processing, Feb. 2024,
doi: 10.48550/arXiv.2306.00442*

Abstract—We present a fast update rule for variational block-sparse Bayesian learning (SBL) methods. Based on a variational Bayesian approximation, we show that iterative updates of probability density functions (PDFs) of the prior precisions and weights can be expressed as a nonlinear first-order recurrence from one estimate of the parameters of the proxy PDFs to the next. In particular, for commonly used prior PDFs such as Jeffrey’s prior, the recurrence relation turns out to be a strictly increasing rational function. This property is the basis for two important analytical results. First, the determination of fixed points by solving for the roots of a polynomial. Second, the determination of the limit of the prior precision after an infinite sequence of update steps. These results are combined into a simplified single-step check for convergence/divergence of each prior precision. Consequently, our proposed criterion significantly reduces the computational complexity of the variational block-SBL algorithm, leading to a remarkable two orders of magnitude improvement in convergence speed shown by simulations. Moreover, the criterion provides valuable insights into the sparsity of the estimators obtained by different prior choices.

This research was partly funded by the Austrian Research Promotion Agency (FFG) within the project SEAMAL Front (project number: 880598). Furthermore, the financial support by the Christian Doppler Research Association, the Austrian Federal Ministry for Digital and Economic Affairs and the National Foundation for Research, Technology and Development is gratefully acknowledged.

D.1 Introduction

Sparse signal models have gained widespread adoption over the last 20 years through the advent of compressed sensing [1]. Generally, the aim of sparse signal reconstruction algorithms is to reconstruct a signal based on a large dictionary matrix as a weighted linear combination of only a few (non-zero) entries of the dictionary. One approach to solve this problem is sparse Bayesian learning (SBL). SBL uses the normal variance mixture model [2] which models the prior precision, i.e., inverse variance, of each weight as a hyperparameter.¹ These hyperparameters are estimated from the data [3–6]. Extensions to block-sparse signal models have been developed, where the weights are grouped into blocks, with only a small number of blocks differing from the zero vector [7–10].

A major shortcoming of the SBL framework is the slow convergence rate of the hyperparameter estimation. However, this problem was alleviated by the introduction of a fast method to maximize the marginal likelihood [4, 5]. In [11], this fast method for maximizing the marginalized likelihood is extended to multiple measurement vectors, which is closely related to block sparsity. Additionally, [11] shows that the hyperparameters can be analytically determined by solving for the roots of a polynomial. However, solving the polynomial with a considerable order was not computationally efficient for the application outlined in [11]. Therefore, an iterative maximization procedure was utilized. In [12], the authors introduced an adaption of the fast marginal likelihood maximization method from [11] tailored to block-sparse models.

Alternatively, SBL can be derived within a variational Bayesian framework [13–16]. In this framework, (approximate) posterior distributions of the parameters are estimated by maximizing the evidence lower bound (ELBO) [17] [18, Ch. 10]. Since posterior probability density functions (PDFs) are estimated instead of point estimates only, the variational approach provides a more complete characterization of the parameters of interest. Furthermore, different hyperpriors, i.e. prior PDFs on the hyperparameters, lead to different estimators. By choosing a parameterized PDF as hyperprior, which encompasses many commonly used hyperpriors as special cases of its parameters, we compare different variants of SBL within the same framework [15, 19]. A fast update rule for the hyperparameters is of critical importance for this analysis. For non-blocked variational SBL such a fast update rule is proposed in [13] by identifying the fixed points of a recurrent sequence of variational updates. However, extending the solution in [13] to block-sparse models is not trivial.

Contribution

The main contribution of this work is the expansion of the variational SBL method, as introduced in [13], to block-sparse models. In particular, we extend and generalize [13, Theorem 1] to block-sparse models. Since [13, Theorem 1] merely provides the condition for one (locally stable) fixed point to be existent, it cannot be used in its original form. We show (i) that for block-sparse models potentially multiple such fixed points exists and (ii) that the fixed points of the variational update sequence

¹Note that SBL can be derived equivalently by using a hierarchical model on either the prior variances or the prior precisions of the weights.

for block-sparse models can be found as the roots of a polynomial. To make the fast update rule [13] applicable to block-sparse models, we need to identify to which fixed point the variational update sequence converges.

Our new theorem, stated as Theorem 1, enables the calculation of the limit of the variational update sequence. I.e., it allows to determine to which fixed point the variational update sequence converges, which is crucial to derive a block-sparse fast update rule. Note that [13, Theorem 1] is regarded as a special case of Theorem 1, where each block consists of only a single component.

Based on Theorem 1, we derive a fast variational block-sparse SBL algorithm with the following advantages.

- The runtime of our algorithm is up to two orders of magnitude smaller than that of similar block-SBL algorithms [9], while simultaneously achieving a better performance in terms of the normalized mean squared error (NMSE) and the probability of correctly recovering the support.
- We can obtain many commonly used hyperprior distributions by modeling the hyperprior as a generalized inverse Gaussian distribution. Thus, the estimators obtained by using different hyperpriors can be analyzed and compared directly using our solution.
- We show that for a certain parameterization of the generalized inverse Gaussian distribution, our proposed method is equivalent to the method of fast marginal likelihood maximization for block-SBL [12]. However, since different parameter settings of the hyperprior lead to different estimators our method is more general.
- We apply our proposed block-SBL to a direction of arrival (DOA) estimation problem and show the benefits compared to the multi-snapshot SBL-based DOA estimation algorithm in [20].

Since a fast block-SBL was already presented by [12], we would like to highlight the similarities and differences between our approach and [12]. While the approach of [12] is based on maximizing the marginalized likelihood similar to [5], we address the issue through the framework of variational Bayesian inference. The variational Bayesian approach allows to consider and analyze different hyperpriors and, thus, different estimators whereas [12] is shown to be equivalent to a special case of our parameterized hyperprior. Another advantage of the variational Bayesian approach is that it can be easily extended to include the estimation of additional parameters such as using a parametrized (infinite) dictionary matrix similar to [14, 21]. Note that the variational Bayesian approach can be easily incorporated into message passing algorithms, such as belief propagation [22]. This integration can be utilized in joint channel estimation and decoding of MIMO communication systems [23–25].

D.2 Overview of the Block-Sparse SBL Framework

D.2.1 System Model

We aim to estimate the weights \mathbf{x} in the linear model

$$\mathbf{y} = \Phi \mathbf{x} + \mathbf{v} \quad (\text{D.1})$$

where \mathbf{y} is the observed signal vector of length N , Φ is an $N \times M$ dictionary matrix and \mathbf{v} is additive white Gaussian noise (AWGN) with precision $\lambda \in \mathbb{R}_{\geq 0}$. Thus, the likelihood function is given by the PDF $p(\mathbf{y}|\mathbf{x}, \lambda) = \mathcal{N}(\mathbf{y}; \Phi\mathbf{x}, \lambda^{-1}\mathbf{I})$.² To allow for both a real-valued and complex-valued signal model, we use a likelihood function that is parameterized by $\rho = \frac{1}{2}$ for the real valued case and $\rho = 1$ for the complex valued case.

We assume that \mathbf{x} is block-sparse, meaning that the M -length vector \mathbf{x} is partitioned into K blocks \mathbf{x}_i of known size d_i each, e.g.

$$\mathbf{x} = \underbrace{[x_1 \ x_2 \ \cdots \ x_{d_1}]}_{\mathbf{x}_1^T} \cdots \cdots \underbrace{[x_{M-d_K+1} \ \cdots \ x_M]}_{\mathbf{x}_K^T}^T \quad (\text{D.2})$$

for which $\mathbf{x}_i = \mathbf{0}$ for all except a few blocks i . Note that we can transform any problem with joint sparsity in groups of non-neighboring weights to the blocked form illustrated in (D.2) by rearranging the columns of the dictionary matrix Φ and the entries of \mathbf{x} . Furthermore, consider the multiple measurement vector (MMV) signal model

$$\mathbf{Y} = \Psi \mathbf{X} + \mathbf{V} \quad (\text{D.3})$$

where $\mathbf{Y} = [\mathbf{y}_1 \ \mathbf{y}_2 \ \cdots \ \mathbf{y}_J]$ is a matrix of J measurement vectors \mathbf{y}_j , Ψ is some dictionary matrix with corresponding weights \mathbf{X} , and \mathbf{V} is a matrix of additive noise. A common assumption is that the sparsity profile is the same for each column in \mathbf{X} , i.e. that all elements in each row of \mathbf{X} are either jointly zero or nonzero. Let $\mathbf{y} = \text{vec}(\mathbf{Y}^T)$, $\mathbf{x} = \text{vec}(\mathbf{X}^T)$, and $\mathbf{v} = \text{vec}(\mathbf{V}^T)$, where $\text{vec}(\mathbf{X})$ denote the operation of stacking all the columns of the $N \times M$ matrix \mathbf{X} into an $NM \times 1$ column vector. Furthermore, let $\Phi = \Psi \otimes \mathbf{I}_J$ be the Kronecker product of Ψ with the $J \times J$ identity matrix \mathbf{I}_J , finding the row-sparse matrix \mathbf{X} in (D.3) is equivalent to finding the block-sparse weight vector \mathbf{x} in (D.1).

The block-sparse or MMV problem has many applications. E.g. the estimation of block-sparse channels in multiple-input multiple-output communication systems [23, 24] or DOA estimation using the MMV model [20]. We provide an example of how to apply the system model in (D.3) to MMV-DOA estimation in Section D.7.

D.2.2 SBL Probabilistic Model

SBL solves the sparse signal reconstruction task of (D.1) through the mechanism of automatic relevance detection [3]. We use a normal variance mixture model [2] and model the prior for each block \mathbf{x}_i as a zero-mean Gaussian PDF

$$p(\mathbf{x}_i|\gamma_i) = \mathcal{N}(\mathbf{x}_i; \mathbf{0}, (\gamma_i \mathbf{B}_i)^{-1}). \quad (\text{D.4})$$

with precision matrix $\gamma_i \mathbf{B}_i$, where \mathbf{B}_i is a known matrix characterizing the intra-block correlation and $\gamma_i \in \mathbb{R}_{\geq 0}$ is a hyperparameter which scales the prior precision matrix of each block [4–6]. The different blocks of weights \mathbf{x}_i are assumed independent, i.e. $p(\mathbf{x}|\boldsymbol{\gamma}) = \prod_{i=1}^K p(\mathbf{x}_i|\gamma_i)$, denoting $\boldsymbol{\gamma} = [\gamma_1 \ \gamma_2 \ \cdots \ \gamma_K]^T$. By estimating $\boldsymbol{\gamma}$ from the data, the relevance of each block is automatically determined. For many blocks i the estimate of γ_i diverges. These blocks are effectively removed from the model as the corresponding estimate of \mathbf{x}_i approaches $\mathbf{0}$.

²We denote the multivariate Gaussian PDF of the variable \mathbf{x} with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$ as $\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) = |\frac{\pi}{\rho} \boldsymbol{\Sigma}|^{-\rho} \exp\{-\rho(\mathbf{x} - \boldsymbol{\mu})^H \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})\}$, where $|\cdot|$ denotes the matrix determinant.

To perform (approximate) Bayesian inference, we place independent hyperprior PDFs $p(\boldsymbol{\gamma}) = \prod_{i=1}^K p(\gamma_i)$ over the unknown hyperparameters γ_i . Given $p(\mathbf{x}|\boldsymbol{\gamma})$ and $p(\boldsymbol{\gamma})$, the prior on the weights is obtained as

$$p(\mathbf{x}) = \int p(\mathbf{x}|\boldsymbol{\gamma})p(\boldsymbol{\gamma}) d\boldsymbol{\gamma}. \quad (\text{D.5})$$

Hence, by specifying different hyperprior PDFs $p(\boldsymbol{\gamma})$, we obtain different priors $p(\mathbf{x})$ and, thus, different estimators. Commonly used PDFs for $p(\gamma_i)$ include the Gamma and inverse Gamma PDFs as well as Jeffrey's prior $p(\gamma_i) \propto \gamma_i^{-1}$ [13, 15, 19]. The generalized inverse Gaussian distribution [26]

$$p(\gamma_i; a_i, b_i, c_i) = \frac{\left(\frac{a_i}{b_i}\right)^{\frac{c_i}{2}} \gamma_i^{c_i-1}}{2K_{c_i}(\sqrt{a_i b_i})} \exp\left\{-\frac{1}{2}(a_i \gamma_i + b_i \gamma_i^{-1})\right\} \quad (\text{D.6})$$

where $K_c(\cdot)$ denotes the modified Bessel function of the second kind, includes the aforementioned PDFs as special cases of the parameters a_i , b_i and c_i as shown in the first and second column of Table D.1. Moreover, this choice of prior for $p(\boldsymbol{\gamma})$ is conjugate for Gaussian $p(\mathbf{x}|\boldsymbol{\gamma})$, which allows to solve the variational updates analytically [15]. Therefore, we consider a generalized inverse Gaussian distribution as hyperprior $p(\gamma_i)$ throughout this work.

To complete the Bayesian model, we also need a prior PDF for the noise precision λ . We use a Gamma PDF with shape ϵ and rate η

$$p(\lambda; \epsilon, \eta) = \frac{\eta^\epsilon}{\Gamma(\epsilon)} \lambda^{\epsilon-1} \exp\{-\eta\lambda\} \quad (\text{D.7})$$

where $\Gamma(\cdot)$ denotes the Gamma function, since this is a conjugate prior for the precision of a Gaussian distribution.

D.2.3 Variational Bayesian Inference

We apply variational inference [17] [18, Ch. 10] to approximate the posterior PDF of the parameters \mathbf{x} , $\boldsymbol{\gamma}$ and λ

$$p(\mathbf{x}, \boldsymbol{\gamma}, \lambda | \mathbf{y}) \propto p(\mathbf{y} | \mathbf{x}, \lambda) p(\mathbf{x} | \boldsymbol{\gamma}) p(\boldsymbol{\gamma}) p(\lambda) \quad (\text{D.8})$$

with a simpler, factorized distribution $q(\mathbf{x}, \boldsymbol{\gamma}, \lambda)$. Specifically, we apply the structured mean-field assumption $q(\mathbf{x}, \boldsymbol{\gamma}, \lambda) = q_{\mathbf{x}}(\mathbf{x}) q_{\lambda}(\lambda) \prod_{i=1}^K q_{\gamma_i}(\gamma_i)$ to approximate the posterior PDF as

$$p(\mathbf{x}, \boldsymbol{\gamma}, \lambda | \mathbf{y}) \approx q_{\mathbf{x}}(\mathbf{x}) q_{\lambda}(\lambda) \prod_{i=1}^K q_{\gamma_i}(\gamma_i). \quad (\text{D.9})$$

Each factor $q_j \in \mathcal{Q} = \{q_{\mathbf{x}}, q_{\lambda}, q_{\gamma_1}, q_{\gamma_2}, \dots, q_{\gamma_K}\}$ of $q(\mathbf{x}, \boldsymbol{\gamma}, \lambda)$ in (D.9) is found by maximizing the ELBO

$$\mathcal{L}(q) = \left\langle \ln \frac{p(\mathbf{x}, \boldsymbol{\gamma}, \lambda | \mathbf{y})}{q(\mathbf{x}, \boldsymbol{\gamma}, \lambda)} \right\rangle_{q(\mathbf{x}, \boldsymbol{\gamma}, \lambda)} \quad (\text{D.10})$$

where $\langle f(\cdot) \rangle_{q(\cdot)}$ denotes the expectation of the function $f(\cdot)$ depending on random variables distributed with PDF $q(\cdot)$.³ Maximizing the ELBO (D.10) is equivalent to minimizing the Kullback-Leibler (KL)-divergence between the true posterior and the approximation q [17] [18, Ch. 10]. The factors $q_j \in \mathcal{Q}$ are updated iteratively using

$$q_j \propto \exp \left\{ \left\langle \ln p(\mathbf{x}, \boldsymbol{\gamma}, \lambda | \mathbf{y}) \right\rangle_{q_{\bar{j}}} \right\} \quad (\text{D.11})$$

while keeping the remaining factors $q_l \in \mathcal{Q} \setminus \{q_j\}$ fixed. We denote $q_{\bar{j}} = \prod_{q_l \in \mathcal{Q} \setminus \{q_j\}} q_l$ as the product of all factors in \mathcal{Q} except for q_j . Inserting the posterior (D.8) into (D.11), we obtain

$$q_{\mathbf{x}}(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \hat{\mathbf{x}}, \hat{\boldsymbol{\Sigma}}) \quad (\text{D.12})$$

as a Gaussian PDF with mean

$$\hat{\mathbf{x}} = \hat{\lambda} \hat{\boldsymbol{\Sigma}} \boldsymbol{\Phi}^H \mathbf{y} \quad (\text{D.13})$$

and covariance matrix

$$\hat{\boldsymbol{\Sigma}} = (\hat{\lambda} \boldsymbol{\Phi}^H \boldsymbol{\Phi} + \mathbf{B}_{\boldsymbol{\gamma}})^{-1} \quad (\text{D.14})$$

where $\hat{\lambda} = \langle \lambda \rangle_{q_{\lambda}}$ is the current estimate of the noise precision, $\mathbf{B}_{\boldsymbol{\gamma}}$ is a block-diagonal matrix with the principal diagonal blocks $\hat{\gamma}_i \mathbf{B}_i$, and $\hat{\gamma}_i = \langle \gamma_i \rangle_{q_{\gamma_i}}$ is the mean of q_{γ_i} [15]. Similarly, by inserting (D.8) into (D.11) we get

$$q_{\gamma_i}(\gamma_i) \propto \gamma_i^{\hat{c}_i - 1} \exp \left\{ -\frac{1}{2}(\hat{a}_i \gamma_i + b_i \gamma_i^{-1}) \right\} \quad (\text{D.15})$$

for $i \in \{1, 2, \dots, K\}$. I.e. q_{γ_i} is the PDF of a generalized inverse Gaussians distribution with parameters

$$\hat{a}_i = 2\rho \langle \mathbf{x}_i^H \mathbf{B}_i \mathbf{x}_i \rangle_{q_{\mathbf{x}}} + a_i \quad (\text{D.16})$$

b_i , and $\hat{c}_i = c_i + \rho d_i$ [15]. For the update of the other factors in \mathcal{Q} only the expectations $\hat{\gamma}_i = \langle \gamma_i \rangle_{q_{\gamma_i}}$ are required, which can be computed from \hat{a}_i , b_i , and \hat{c}_i as [15, 26]

$$\hat{\gamma}_i = \langle \gamma_i \rangle_{q_{\gamma_i}} = \frac{\sqrt{b_i} K_{\hat{c}_i+1}(\sqrt{\hat{a}_i b_i})}{\sqrt{\hat{a}_i} K_{\hat{c}_i}(\sqrt{\hat{a}_i b_i})}. \quad (\text{D.17})$$

Simplified expressions of the update equation (D.17) for different selections of the parameters settings a_i , b_i and c_i are listed in the third column of Table D.1. Note, that if Jeffery's prior is used, the simplified form $\hat{\gamma}_i = d_i / \langle \mathbf{x}_i^H \mathbf{B}_i \mathbf{x}_i \rangle_{q_{\mathbf{x}}}$ of (D.17) given in the last row of Table D.1 is equivalent to the update rule [9, Eq. (4)] obtained by applying the expectation-maximization (EM)-algorithm.⁴ Hence, the presented fast solution can also be applied to [9] and similar EM-based block-SBL algorithms.

³In the following we use q as a shorthand for $q(\mathbf{x}, \boldsymbol{\gamma}, \lambda)$ and similarly q_j as a shorthand for $q_j(j)$ with $j \in \{\mathbf{x}, \lambda, \gamma_1, \gamma_2, \dots, \gamma_K\}$.

⁴Reference [9] uses prior variances in their model instead of the prior precisions we use throughout the paper. Thus, the variable $\hat{\gamma}_i$ in our derivations corresponds to γ_i^{-1} in [9].

Table D.1: Fast update rules of $\hat{\gamma}_i$ for different parametrizations of $p(\gamma_i)$, expanding on [15, Table 1]

Hyperprior Shape	Hyperprior Parameters	Update of $\hat{\gamma}_i$	$f_i(\hat{\gamma}_i)$	Fast Update Polynomial $G_i(\hat{\gamma}_i)$	Sparse
Generalized inverse Gaussian	$a_i > 0, b_i > 0, c_i \in \mathbb{R}$	$\sqrt{\frac{b_i}{\hat{a}_i} \frac{K_{\hat{c}_i+1}(\sqrt{\hat{a}_i} b_i)}{K_{\hat{c}_i}(\sqrt{\hat{a}_i} b_i)}}$	-	-	-
Inverse Gamma	$a_i \rightarrow 0, b_i > 0, c_i = -\rho d_i - \frac{1}{2}$	$\sqrt{\frac{b_i}{2\rho \langle \mathbf{x}_i^H \mathbf{B}_i \mathbf{x}_i \rangle_{q\mathbf{x}}}}$	$\sqrt{\frac{b_i A_i(\hat{\gamma}_i)}{2\rho B_i(\hat{\gamma}_i)}}$	$b_i A_i(\hat{\gamma}_i) - 2\rho \hat{\gamma}_i^2 B_i(\hat{\gamma}_i)$	No
Gamma	$a_i > 0, b_i \rightarrow 0, c_i > 0$	$\frac{c_i + \rho d_i}{\rho \langle \mathbf{x}_i^H \mathbf{B}_i \mathbf{x}_i \rangle_{q\mathbf{x}} + \frac{a_i}{2}}$	$\frac{c_i + \rho d_i}{\rho \frac{B_i(\hat{\gamma}_i)}{A_i(\hat{\gamma}_i)} + \frac{a_i}{2}}$	$(c_i + \rho d_i) A_i(\hat{\gamma}_i) - \hat{\gamma}_i (\rho B_i(\hat{\gamma}_i) + \frac{a_i}{2} A_i(\hat{\gamma}_i))$	No
Scaled Jeffrey's $p(\gamma_i) \propto \gamma_i^{-1+c_i}$	$a_i \rightarrow 0, b_i \rightarrow 0, c_i > -\rho d_i$	$\frac{c_i + \rho d_i}{\rho \langle \mathbf{x}_i^H \mathbf{B}_i \mathbf{x}_i \rangle_{q\mathbf{x}}}$	$\frac{c_i + \rho d_i}{\rho \frac{B_i(\hat{\gamma}_i)}{A_i(\hat{\gamma}_i)}}$	$(c_i + \rho d_i) A_i(\hat{\gamma}_i) - \rho \hat{\gamma}_i B_i(\hat{\gamma}_i)$	Yes if $c_i > 0$
Jeffrey's $p(\gamma_i) \propto \gamma_i^{-1}$	$a_i \rightarrow 0, b_i \rightarrow 0, c_i \rightarrow 0$	$\frac{d_i}{\langle \mathbf{x}_i^H \mathbf{B}_i \mathbf{x}_i \rangle_{q\mathbf{x}}}$	$\frac{d_i A_i(\hat{\gamma}_i)}{B_i(\hat{\gamma}_i)}$	$d_i A_i(\hat{\gamma}_i) - \hat{\gamma}_i B_i(\hat{\gamma}_i)$	Yes

Note that our model is slightly different from the one used in [15]. The hyperparameters γ_i correspond to z_i^{-1} in [15]. This reparametrization results in a swap of the hyperprior parameters a_i and b_i , and a change in the sign of c_i compared to [15, Table 1]. Furthermore, the Gamma hyperprior PDF corresponds to the inverse Gamma mixing distribution of [15, Table 1] and vice versa.

Finally, by inserting (D.8) into (D.11) we obtain

$$q_\lambda(\lambda) \propto \lambda^{\hat{\epsilon}-1} \exp\{-\hat{\eta}\lambda\} \quad (\text{D.18})$$

i.e. a Gamma PDF with shape $\hat{\epsilon} = \rho N + \epsilon$ and rate $\hat{\eta} = \rho(\|\mathbf{y} - \Phi \hat{\mathbf{x}}\|^2 + \text{tr}(\Phi^H \Phi \hat{\Sigma})) + \eta$. The expectation $\hat{\lambda}$ of q_λ is

$$\hat{\lambda} = \langle \lambda \rangle_{q_\lambda} = \frac{\rho N + \epsilon}{\rho(\|\mathbf{y} - \Phi \hat{\mathbf{x}}\|^2 + \text{tr}(\Phi^H \Phi \hat{\Sigma})) + \eta}. \quad (\text{D.19})$$

D.3 Variational Fast Solution

D.3.1 Derivation of the Fast Solution

In the variational Bayesian framework, a solution is typically obtained by iterating the update equations (D.13), (D.14), (D.17) and (D.19) until convergence. However, the convergence of the hyperparameters $\hat{\gamma}_i$ can be slow due to the cyclic dependency of $\hat{\gamma}_i$ on $\hat{\mathbf{x}}$ and $\hat{\Sigma}$ and vice versa. To accelerate convergence, we analyse how the parameters $\hat{\gamma}_i$ of the distribution q_{γ_i} of a single hyperprior behave if updates of $q_{\mathbf{x}}$ and q_{γ_i} are repeated ad infinitum. Following the approach of [13], we consider the sequence of estimates $\{\hat{\gamma}_i^{[n]}\}_{n=1}^\infty$ obtained by repeated cycles of updating $q_{\mathbf{x}}$ followed by updating q_{γ_i} . For specific parameter settings a_i , b_i and c_i , each element of the sequence can be calculated by the simplified version of (D.17) listed in the third column of Table D.1. We are interested in knowing if the sequence converges and, if it does converge, in its limit.

Note, that the expression $h_i(\hat{\gamma}_i) \triangleq \langle \mathbf{x}_i^H \mathbf{B}_i \mathbf{x}_i \rangle_{q_{\mathbf{x}}}$ in the update equations for $\hat{\gamma}_i$ given in the third column of Table D.1 is a function of the previous estimate of $\hat{\gamma}_i$ through (D.13) and (D.14). Hence, the sequence of estimates $\hat{\gamma}_i^{[n]}$ obtained by repeatedly updating $q_{\mathbf{x}}$ followed by q_{γ_i} can be viewed as generated by a first-order recurrence $f_i : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ which maps from one estimate in the sequence to the next as $\hat{\gamma}_i^{[n+1]} = f_i(\hat{\gamma}_i^{[n]})$. If the update sequence converges, the limit must be a fixed point $\hat{\gamma}_i^*$ of this recurrent relation that fulfills

$$f(\hat{\gamma}_i^*) - \hat{\gamma}_i^* = 0. \quad (\text{D.20})$$

As we show in Section D.3.2, the function $h_i(\hat{\gamma}_i)$ can be expressed as a rational function

$$h_i(\hat{\gamma}_i) \triangleq \langle \mathbf{x}_i^H \mathbf{B}_i \mathbf{x}_i \rangle_{q_{\mathbf{x}}} = \frac{B_i(\hat{\gamma}_i)}{A_i(\hat{\gamma}_i)} \quad (\text{D.21})$$

where $A_i(\hat{\gamma}_i)$ and $B_i(\hat{\gamma}_i)$ are two polynomials in $\hat{\gamma}_i$. We insert the simplified update equations f_i , given in the fourth column of Table D.1, and (D.21) into (D.20) to find the fixed points $\hat{\gamma}_i^*$ as solutions to the polynomial equation $G_i(\hat{\gamma}_i) = 0$ after a few algebraic manipulations. Thus, the sets of fixed points can be obtained as

$$\mathcal{G}_i \triangleq \{\hat{\gamma}_i \in \mathbb{R}_{\geq 0} : G_i(\hat{\gamma}_i) = 0\}. \quad (\text{D.22})$$

Different polynomials $G_i(\hat{\gamma}_i)$ are obtained for the different shapes of the hyperprior $p(\gamma_i)$ listed in the first column of Table D.1. For the readers convenience these polynomials are included in the fifth column of Table D.1 as well.

Moreover, Appendix D.9.1 shows that $h_i(\hat{\gamma}_i) = B_i(\hat{\gamma}_i)/A_i(\hat{\gamma}_i)$ is strictly decreasing. It follows, that the recurrence relations f_i given in the fourth column of Table D.1 are smooth and strictly increasing functions since they are the reciprocal of a smooth and strictly decreasing function $h_i(\hat{\gamma}_i) = B_i(\hat{\gamma}_i)/A_i(\hat{\gamma}_i)$. Thus, the following variant of the monotone convergence theorem [27, Theorem 3.14] proves under which condition the sequence $\{\hat{\gamma}_i^{[n]}\}_{n=1}^\infty$ converges or diverges based on the set of fixed points \mathcal{G}_i .

Theorem 1. *For any initial condition $\hat{\gamma}_i^{[0]} \geq 0$, the convergence of the sequence of estimates $\{\hat{\gamma}_i^{[n]}\}_{n=1}^\infty$ generated as $\hat{\gamma}_i^{[n+1]} = f_i(\hat{\gamma}_i^{[n]})$ by any strictly increasing first-order recurrence $f_i : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, such as the ones listed in the fourth column of Table D.1, is governed by the sets of fixed points $\mathcal{G}_i^+ = \{\hat{\gamma}_i \in (\hat{\gamma}_i^{[0]}, \infty) : f_i(\hat{\gamma}_i) - \hat{\gamma}_i = 0\}$ greater than $\hat{\gamma}_i^{[0]}$, and fixed points $\mathcal{G}_i^- = \{\hat{\gamma}_i \in [0, \hat{\gamma}_i^{[0]}] : f_i(\hat{\gamma}_i) - \hat{\gamma}_i = 0\}$ smaller than or equal to $\hat{\gamma}_i^{[0]}$ as*

$$\lim_{n \rightarrow \infty} \hat{\gamma}_i^{[n]} = \begin{cases} \infty & \text{if } f_i(\hat{\gamma}_i^{[0]}) > \hat{\gamma}_i^{[0]} \text{ and } \mathcal{G}_i^+ = \emptyset \\ \min \mathcal{G}_i^+ & \text{if } f_i(\hat{\gamma}_i^{[0]}) > \hat{\gamma}_i^{[0]} \text{ and } \mathcal{G}_i^+ \neq \emptyset \\ \max \mathcal{G}_i^- & \text{if } f_i(\hat{\gamma}_i^{[0]}) \leq \hat{\gamma}_i^{[0]} \end{cases} \quad (\text{D.23})$$

where \emptyset denotes the empty set.

Proof. Since f_i is strictly increasing, the sequence $\{\hat{\gamma}_i^{[n]}\}_{n=1}^\infty$ is either strictly increasing if $f_i(\hat{\gamma}_i^{[0]}) > \hat{\gamma}_i^{[0]}$, or strictly decreasing if $f_i(\hat{\gamma}_i^{[0]}) < \hat{\gamma}_i^{[0]}$, while the case $f_i(\hat{\gamma}_i^{[0]}) = \hat{\gamma}_i^{[0]}$ is trivial. By definition every fixed point $\hat{\gamma}_i^* \in \mathcal{G}_i = \mathcal{G}_i^- \cup \mathcal{G}_i^+$ must fulfil $f_i(\hat{\gamma}_i^*) = \hat{\gamma}_i^*$. Since f_i is strictly increasing, it follows that the sequence $\{\hat{\gamma}_i^{[n]}\}$ is upper bounded by $\min \mathcal{G}_i^+$ (if such a fixed point exists) and lower bounded by $\max \mathcal{G}_i^-$. The monotone convergence theorem [27, Theorem 3.14] states, that the sequence $\{\hat{\gamma}_i^{[n]}\}_{n=1}^\infty$ converges to the upper bound if it is increasing or to the lower bound if it is decreasing. If the sequence is decreasing, it can not diverge since by definition $f_i(\hat{\gamma}_i) \geq 0$. Thus, a fixed point $\hat{\gamma}_i^*$ must exist in the interval $[0, \hat{\gamma}_i^{[0]}]$ and \mathcal{G}_i^- can not be empty. If the sequence is increasing and no fixed point $\hat{\gamma}_i^*$ exists in the interval $(\hat{\gamma}_i^{[0]}, \infty)$, i.e. if \mathcal{G}_i^+ is empty, we can always find some small constant $\Delta : f_i(\hat{\gamma}_i^{[n]}) - \hat{\gamma}_i^{[n]} \geq \Delta > 0 \forall n \in \mathbb{Z}_{>0}$. Hence, for any arbitrarily large positive constant $C : \hat{\gamma}_i^{[n]} > C \forall n > \lceil \frac{C}{\Delta} \rceil$, i.e. the sequence diverges.⁵ ■

We derive a fast update rule for q_{γ_i} and $q_{\mathbf{x}}$ as follows. First, the set of fixed points \mathcal{G}_i is obtained by solving for the roots of the polynomial $G_i(\hat{\gamma}_i)$ in $\mathbb{R}_{\geq 0}$, e.g. by solving for the eigenvalues of the companion matrix [28].

Once the fixed points \mathcal{G}_i are obtained, we apply Theorem 1 and (D.23) to update our estimate of $\hat{\gamma}_i$ obtained by iterating the variational update sequence $q_{\gamma_i} \rightarrow q_{\mathbf{x}} \rightarrow q_{\gamma_i} \rightarrow q_{\mathbf{x}} \rightarrow \dots$ ad infinitum with $\hat{\gamma}_i = \lim_{n \rightarrow \infty} \hat{\gamma}_i^{[n]}$. Each fast variational update (D.23) is guaranteed to increase the ELBO, since it is the result of an infinite sequence of update steps where each single step increases the ELBO.

⁵We use $\lceil \cdot \rceil$ to denotes the operation of rounding up to the next highest integer number.

Discussion of Theorem 1

While [13, Theorem 1] proofs under which condition a (locally stable) fixed point exists, that theorem is only applicable for the non-blocked case $d_i = 1 \forall i$. In the block-sparse case $d_i > 1$, several such fixed points might exist since the polynomial $G_i(\hat{\gamma}_i)$ can have up to P positive roots with P denoting the order of $G_i(\hat{\gamma}_i)$. To find the limit of the sequence $\{\hat{\gamma}_i^{[n]}\}_{n=1}^\infty$ for $d_i > 1$, we need a method to determine to which of those fixed points the sequence converges. While such a method can not be determined by [13, Theorem 1], (D.23) from our Theorem 1 allows to determine to which of the fixed points the sequence converges.

D.3.2 Derivation of the Update-Polynomials

In this subsection, we show how the expectation $\langle \mathbf{x}_i^H \mathbf{B}_i \mathbf{x}_i \rangle_{q_{\mathbf{x}}}$ in the update of $\hat{\gamma}_i$ can be simplified into a rational function $h_i(\hat{\gamma}_i) = \frac{B_i(\hat{\gamma}_i)}{A_i(\hat{\gamma}_i)}$ as in (D.21). Note, that

$$\langle \mathbf{x}_i^H \mathbf{B}_i \mathbf{x}_i \rangle_{q_{\mathbf{x}}} = \hat{\mathbf{x}}_i^H \mathbf{B}_i \hat{\mathbf{x}}_i + \text{tr}(\mathbf{B}_i \hat{\Sigma}_i) \quad (\text{D.24})$$

is an expectation over a Gaussian PDF $q_{\mathbf{x}}$, where $\hat{\Sigma}_i$ denotes the $d_i \times d_i$ submatrix on the principal diagonal of $\hat{\Sigma}$ which corresponds to the (marginal) covariance of the i -th block \mathbf{x}_i in $q_{\mathbf{x}}$ [29, Eq. (378)]. Let $\mathbf{E}_i = [\mathbf{0} \ \mathbf{I}_{d_i} \ \mathbf{0}]^T$ be an $M \times d_i$ selection matrix such that $\hat{\mathbf{x}}_i = \mathbf{E}_i^T \hat{\mathbf{x}}$ and $\hat{\Sigma}_i = \mathbf{E}_i^T \hat{\Sigma} \mathbf{E}_i$, and let $\mathbf{L}_i \mathbf{L}_i^H = \mathbf{B}_i$ be the Cholesky decomposition of \mathbf{B}_i .⁶ First, we investigate the trace term

$$\text{tr}(\mathbf{B}_i \hat{\Sigma}_i) = \text{tr}(\mathbf{L}_i^H \mathbf{E}_i^T \hat{\Sigma} \mathbf{E}_i \mathbf{L}_i) \quad (\text{D.25})$$

in (D.24). To make the dependence on $\hat{\gamma}_i$ explicit, we write $\hat{\Sigma}$ as

$$\hat{\Sigma} = \left(\lambda \Phi^H \Phi + \sum_{k=1, k \neq i}^K \hat{\gamma}_k \mathbf{E}_k \mathbf{B}_k \mathbf{E}_k^T + \hat{\gamma}_i \mathbf{E}_i \mathbf{B}_i \mathbf{E}_i^T \right)^{-1} \quad (\text{D.26})$$

and apply the Woodbury matrix identity

$$\hat{\Sigma} = \bar{\Sigma} - \bar{\Sigma} \mathbf{E}_i (\hat{\gamma}_i^{-1} \mathbf{B}_i^{-1} + \mathbf{E}_i^T \bar{\Sigma} \mathbf{E}_i)^{-1} \mathbf{E}_i^T \bar{\Sigma} \quad (\text{D.27})$$

where $\bar{\Sigma} = (\lambda \Phi^H \Phi + \sum_{k=1, k \neq i}^K \hat{\gamma}_k \mathbf{E}_k \mathbf{B}_k \mathbf{E}_k^T)^{-1}$. Let $\bar{\Sigma}_i = \mathbf{E}_i^T \bar{\Sigma} \mathbf{E}_i$, and let $\mathbf{U}_i \mathbf{S}_i \mathbf{U}_i^H$ be the eigendecomposition of $\mathbf{L}_i^H \bar{\Sigma}_i \mathbf{L}_i$, such that the columns of \mathbf{U}_i are the eigenvectors of $\mathbf{L}_i^H \bar{\Sigma}_i \mathbf{L}_i$ and \mathbf{S}_i is a diagonal matrix with the corresponding eigenvalues $s_{i,l}$ on its main diagonal and $\mathbf{L}_i^H \bar{\Sigma}_i \mathbf{L}_i = \mathbf{U}_i \mathbf{S}_i \mathbf{U}_i^H$. We insert (D.27) into (D.25) and use the identity

$$(\hat{\gamma}_i^{-1} \mathbf{B}_i^{-1} + \bar{\Sigma}_i)^{-1} = \mathbf{L}_i \mathbf{U}_i (\hat{\gamma}_i^{-1} \mathbf{I}_{d_i} + \mathbf{S}_i)^{-1} \mathbf{U}_i^H \mathbf{L}_i^H \quad (\text{D.28})$$

to rewrite the trace term in (D.24) as

$$\begin{aligned} \text{tr}(\mathbf{B}_i \hat{\Sigma}_i) &= \text{tr}(\mathbf{S}_i) - \text{tr}(\mathbf{S}_i (\hat{\gamma}_i^{-1} \mathbf{I}_{d_i} + \mathbf{S}_i)^{-1} \mathbf{S}_i) \\ &= \sum_{l=1}^{d_i} \frac{s_{i,l}}{1 + \hat{\gamma}_i s_{i,l}}. \end{aligned} \quad (\text{D.29})$$

⁶For the case of $\mathbf{B}_i = \text{diag}(\mathbf{b}_i)$ being a diagonal matrix with the elements of the vector \mathbf{b}_i on its main diagonal, the Cholesky decomposition results in the matrix $\mathbf{L}_i = \text{diag}(\sqrt{\mathbf{b}_i})$. Similarly, if $\mathbf{B}_i = \mathbf{I}$ then it follows that $\mathbf{L}_i = \mathbf{I}$. Subsequently, \mathbf{L}_i can be removed from the definition of $\mathbf{q}_i = \hat{\lambda} \mathbf{U}_i^H \mathbf{E}_i^T \bar{\Sigma} \Phi^H \mathbf{y}$, and s_i are the eigenvalues of $\bar{\Sigma}_i$.

Next, we investigate the expression $\hat{\mathbf{x}}_i^H \mathbf{B}_i \hat{\mathbf{x}}_i$ in (D.24). Using (D.13), (D.27) and (D.28), we write

$$\begin{aligned} \hat{\mathbf{x}}_i^H \mathbf{B}_i \hat{\mathbf{x}}_i &= \hat{\gamma}_i^2 \mathbf{y}^H \Phi \bar{\Sigma} \mathbf{E}_i \mathbf{L}_i \mathbf{U}_i [\mathbf{I}_{d_i} - 2\mathbf{S}_i(\hat{\gamma}_i^{-1} \mathbf{I}_{d_i} + \mathbf{S}_i)^{-1} \\ &\quad + \mathbf{S}_i^2(\hat{\gamma}_i^{-1} \mathbf{I}_{d_i} + \mathbf{S}_i)^{-2}] \mathbf{U}_i^H \mathbf{L}_i^H \mathbf{E}_i^T \bar{\Sigma} \Phi^H \mathbf{y} \end{aligned} \quad (\text{D.30})$$

after applying a few algebraic manipulations. Equation (D.30) is a quadratic form $\mathbf{q}_i^H \mathbf{A}_i \mathbf{q}_i$ of the vector

$$\mathbf{q}_i = \hat{\lambda} \mathbf{U}_i^H \mathbf{L}_i^H \mathbf{E}_i^T \bar{\Sigma} \Phi^H \mathbf{y} \quad (\text{D.31})$$

and a diagonal matrix $\mathbf{A}_i = \mathbf{I}_{d_i} - 2\mathbf{S}_i(\hat{\gamma}_i^{-1} \mathbf{I}_{d_i} + \mathbf{S}_i)^{-1} + \mathbf{S}_i^2(\hat{\gamma}_i^{-1} \mathbf{I}_{d_i} + \mathbf{S}_i)^{-2}$, which can be expressed as

$$\hat{\mathbf{x}}_i^H \mathbf{B}_i \hat{\mathbf{x}}_i = \sum_{l=1}^{d_i} \frac{|q_{i,l}|^2}{(1 + \hat{\gamma}_i s_{i,l})^2} \quad (\text{D.32})$$

where $q_{i,l}$ is the l -th entry of the vector \mathbf{q}_i .

We define

$$A_i(\hat{\gamma}_i) = \prod_{l=1}^{d_i} (1 + \hat{\gamma}_i s_{i,l})^2 \quad (\text{D.33})$$

as a polynomial of order $2d_i$ and

$$B_i(\hat{\gamma}_i) = \sum_{l=1}^{d_i} (\hat{\gamma}_i s_{i,l}^2 + |q_{i,l}|^2 + s_{i,l}) \prod_{j=1, j \neq i}^{d_i} (1 + \hat{\gamma}_i s_{i,j})^2 \quad (\text{D.34})$$

as a polynomial of order $2d_i - 1$. We insert (D.29) and (D.32) into (D.24), we arrive at (D.21) after some algebraic manipulations. Thus, proving that the expectation $\langle \mathbf{x}_i^H \mathbf{B}_i \mathbf{x}_i \rangle_{q_{\mathbf{x}}}$ in the update of \hat{a}_i is indeed a rational function.

D.4 Efficient Implementation

The solution of iterating updates of $q_{\mathbf{x}}$ and q_{γ_i} ad infinitum can be obtained by the following steps if the hyperprior PDF is one of the special cases of the generalized inverse Gaussian distribution listed in Table D.1.

1. Calculate $\bar{\Sigma}$ and the eigenvalue decomposition of $\mathbf{L}_i^H \bar{\Sigma}_i \mathbf{L}_i$ to obtain \mathbf{q}_i and \mathbf{S}_i . Furthermore, calculate the coefficients of the polynomials $A_i(\hat{\gamma}_i)$ and $B_i(\hat{\gamma}_i)$ using (D.33) and (D.34), respectively.
2. Solve for the real and positive roots of the polynomial $G_i(\hat{\gamma}_i)$ and construct the set \mathcal{G}_i from (D.22).
3. Apply Theorem 1 to update $\hat{\gamma}_i$ using (D.23).

Note, that all rows and columns of $\hat{\Sigma}$ which correspond to a block i with $\hat{\gamma}_i = \infty$ will be zero. Hence, the algorithm can be implemented efficiently by considering

Algorithm 1 Fast Variational Block-Sparse Bayesian Learning**Input:** Observations \mathbf{y} , dictionary Φ , precision matrices \mathbf{B}_i .**Output:** Weights $\hat{\mathbf{x}}$, hyperparameters γ , noise precision $\hat{\lambda}$.Initialize $n = 1$, $\hat{\lambda} = \frac{2N}{\|\mathbf{y}\|^2}$ and all $\hat{\gamma}_i = \infty$.**while** not converged **do** **for all** blocks $i \in \{1, 2, \dots, K\}$ **do** $\bar{\Sigma} \leftarrow (\hat{\lambda} \Phi^H \Phi + \sum_{k=1, k \neq i}^K \hat{\gamma}_k \mathbf{E}_k \mathbf{B}_k \mathbf{E}_k^T)^{-1}$. Calculate \mathbf{U}_i , \mathbf{S}_i and \mathbf{q}_i . Obtain the roots of $G_i(\hat{\gamma}_i)$ in $\mathbb{R}_{\geq 0}$. **if** $n \leq 3$ **then** Update $\hat{\gamma}_i$ using (D.23), assuming $\hat{\gamma}_i^{[0]} = 0$. **else** Update $\hat{\gamma}_i$ using (D.23). **end if** **end for** $\hat{\Sigma} \leftarrow (\hat{\lambda} \Phi^H \Phi + \mathbf{B}_\gamma)^{-1}$. $\hat{\mathbf{x}} \leftarrow \hat{\lambda} \hat{\Sigma} \Phi^H \mathbf{y}$. Update $\hat{\lambda}$ using (D.19). $n \leftarrow n + 1$.**end while**

only the nonzero rows and columns of $\hat{\Sigma}$. As detailed in Algorithm 1,⁷ we initialize the algorithm with $\hat{\lambda} = 2N/\|\mathbf{y}\|^2$ and $\hat{\gamma}_i = \infty \forall i$ (i.e. with an empty model).

We cycle through all blocks $i \in \{1, 2, \dots, K\}$ to update $\hat{\gamma}_i$ using our fast solution (D.23). We perform an update of the noise precision $\hat{\lambda}$ after we updated all $\hat{\gamma}_i$. These steps are iterated until no blocks have been added or removed from the model during the last iteration and the change in the ELBO from one iteration to the next is below a threshold. To avoid the trivial maximum of the ELBO at $\hat{\gamma}_i = \infty \forall i$ in case scaled Jeffrey's prior is used as hyperprior discussed in Section D.5.2, we update $\hat{\gamma}_i$ as if the previous update was $\hat{\gamma}_i^{[0]} = 0$ in the first three iterations. Note, that these modified updates of $\hat{\gamma}_i$ might decrease the ELBO. Nevertheless, we perform such modified updates only in the first three iteration to achieve a better initialization. After the third iteration we only perform unmodified updates, which are guaranteed to increase the ELBO. Therefore, Algorithm 1 is guaranteed to converge to a (local) maximum of the ELBO.

Computational Complexity

Let $\hat{M} = \|\hat{\mathbf{x}}\|_0$ denote the number of nonzero elements of $\hat{\mathbf{x}}$, the computational complexity for each step is as follows:

- The calculation of $\bar{\Sigma}$ is of complexity $\mathcal{O}(\hat{M}^3)$.
- The eigenvalue decomposition of $\mathbf{L}_i^H \bar{\Sigma}_i \mathbf{L}_i$ is of complexity $\mathcal{O}(d_i^3)$.

⁷Matlab code for the proposed algorithm is available at <https://gitlab.com/jmoeder1/fast-variational-block-sparse-bayesian-learning>

- Calculating the coefficients of the polynomials $A_i(\hat{\gamma}_i)$ and $B_i(\hat{\gamma}_i)$ is of complexity $\mathcal{O}(d_i^3)$.
- Solving for the roots of the polynomial $G_i(\hat{\gamma}_i) = 0$ is of complexity $\mathcal{O}(d_i^3)$.
- Updating $\hat{\gamma}_i$ according to Theorem 1 with \mathcal{G}_i already obtained is of complexity $\mathcal{O}(d_i)$.

Assuming that d_i is a small constant with respect to \hat{M} , the by far most computationally complex operation is the computation of $\hat{\Sigma}$ with complexity $\mathcal{O}(\hat{M}^3)$. Applying the EM-algorithm directly as in [9], the computationally most demanding operation is again the matrix inversion required for the covariance matrix, which is of complexity $\mathcal{O}(N^3)$. It follows from the sparsity assumption that $\hat{M} \ll N$ and, thus, our algorithm is of lower computational complexity than the ones proposed in [9]. This difference is further enhanced by the fact, that we check for convergence of $\hat{\gamma}_i$ in a single step. Hence, only a few iterations are needed, as demonstrated in Section D.6.

D.5 Analysis of the Algorithm

D.5.1 Sparsity Analysis

To perform sparse signal reconstruction, we are interested in hyperpriors $p(\gamma_i)$ that result in a sparse estimate $\hat{\mathbf{x}}$ when applying Algorithm 1. However, not all hyperpriors $p(\gamma_i)$ result in a sparse estimate [15]. A sparse estimate is obtained if the mean $\hat{\gamma}_i$ of the estimated proxy PDF q_{γ_i} diverges for many blocks i . Since $\hat{\gamma}_i$ is obtained as the limit $n \rightarrow \infty$ of the sequence $\hat{\gamma}_i^{[n+1]} = f_i(\hat{\gamma}_i^{[n]})$, a necessary condition for $\hat{\gamma}_i$ to diverge is that applying the recurrent map $\hat{\gamma}_i^{[n+1]} = f_i(\hat{\gamma}_i^{[n]})$ increases the value of $\hat{\gamma}_i^{[n+1]}$ compared to $\hat{\gamma}_i^{[n]}$ by at least some small constant Δ for any values of $\hat{\gamma}_i^{[n]}$ larger than some value γ'_i , i.e. $\exists \gamma'_i, \Delta \in \mathbb{R}_{>0} : \forall \hat{\gamma}_i > \gamma'_i, f_i(\hat{\gamma}_i) - \hat{\gamma}_i > \Delta$.

Using the definitions of the map f_i given in the fourth column of Table D.1, it is readily shown by a few algebraic manipulations that the condition $f_i(\hat{\gamma}_i) - \hat{\gamma}_i > \Delta$ is equivalent to $G_i(\hat{\gamma}_i) > H_i(\hat{\gamma}_i)$, where $G_i(\hat{\gamma}_i)$ is the “fast update polynomial” given in the fifth column of Table D.1, and $H_i(\hat{\gamma}_i)$ is some polynomial of $\hat{\gamma}_i$ that depends on the chosen hyperprior $p(\gamma_i)$. Let the polynomial $G_i(\hat{\gamma}_i)$ be expressed as $G_i(\hat{\gamma}_i) = \sum_{k=0}^{P_G} g_{i,k} \hat{\gamma}_i^k$, where P_G denotes the polynomial order of G_i , and let P_H denote the polynomial order of H_i . In the following, we analyse the expression $G_i(\hat{\gamma}_i) > H_i(\hat{\gamma}_i)$ as $\hat{\gamma}_i \rightarrow \infty$ via the polynomial orders P_G and P_H and the coefficients $g_{i,k}$, specifically, the highest-order coefficient g_{i,P_G} . Using $f_i(\hat{\gamma}_i) - \hat{\gamma}_i > \Delta$, we insert the definitions for $f_i(\hat{\gamma}_i)$ given in Table D.1 as well as (D.33) and (D.34), to arrive at the following for different hyperprior PDFs $p(\gamma)$.

Inverse Gamma PDF

In case of an inverse Gamma PDF as hyperprior, the polynomial $G_i(\hat{\gamma}_i)$ is of order $P_G = 2d_i + 1$ and $H_i(\hat{\gamma}_i) = 2\rho\Delta^2 B_i(\hat{\gamma}_i)$ is of order $P_H = 2d_i - 1$. Since $P_G > P_H$ and the highest-order coefficient $g_{i,P_G} = -2\rho d_i \prod_{l=1}^{d_i} s_{i,l}^2$ is negative, $G_i(\hat{\gamma}_i) > H_i(\hat{\gamma}_i)$ is not fulfilled as $\hat{\gamma}_i \rightarrow \infty$, resulting in a non-sparse estimate $\hat{\mathbf{x}}$.

Gamma PDF

In case of a Gamma PDF as hyperprior, the polynomial $G_i(\hat{\gamma}_i)$ is of order $P_G = 2d_i + 1$ and $H_i(\hat{\gamma}_i) = \Delta(\rho B_i(\hat{\gamma}_i) + \frac{a_i}{2} A(\hat{\gamma}_i))$ is of order $P_H = 2d_i$. Since $P_G > P_H$ and the highest-order coefficient $g_{i,P_G} = -\frac{a_i}{2} \prod_{l=1}^{d_i} s_{i,l}^2$ is strictly negative, $G_i(\hat{\gamma}_i) > H_i(\hat{\gamma}_i)$ is not fulfilled as $\hat{\gamma}_i \rightarrow \infty$, resulting in a non-sparse estimate $\hat{\mathbf{x}}$.

Scaled Jeffery's PDF

For the scaled Jeffrey's PDF, i.e. the improper prior $p(\gamma_i) \propto \gamma_i^{-1+c_i}$, the polynomial $G_i(\hat{\gamma}_i)$ is of order $P_G = 2d_i$ and $H_i(\hat{\gamma}_i) = \rho \Delta B_i(\hat{\gamma}_i)$ is of order $P_H = 2d_i - 1$. In this case again $P_G > P_H$. However, the sign of the highest-order coefficient $g_{i,P_G} = c_i \prod_{l=1}^{d_i} s_{i,l}^2$ is equal to the sign of the parameter c_i . If $c_i < 0$, then $g_{i,P_G} < 0$ and $G_i(\hat{\gamma}_i) > H_i(\hat{\gamma}_i)$ is not fulfilled as $\hat{\gamma}_i \rightarrow \infty$, resulting in a non-sparse estimate $\hat{\mathbf{x}}$. On the other hand, if $c_i > 0$, then $g_{i,P_G} > 0$ and $G_i(\hat{\gamma}_i) > H_i(\hat{\gamma}_i)$ is fulfilled as $\hat{\gamma}_i \rightarrow \infty$, resulting in a (potentially) sparse estimate $\hat{\mathbf{x}}$.

Jeffrey's PDF

Consider Jeffrey's prior as a special case of the scaled Jeffrey's prior with $c_i = 0$. It is readily obtained that the coefficient $g_{i,2d_i} = 0$. Thus, $G_i(\hat{\gamma}_i)$ is of the same order as $H_i(\hat{\gamma}_i)$, $P_G = P_H = 2d_i - 1$. In this case, the condition $G_i(\hat{\gamma}_i) > H_i(\hat{\gamma}_i)$ depends on the values of $s_{i,l}$ and $|q_{i,l}|^2$ in a non-trivial manner, resulting in a (potentially) sparse estimate $\hat{\mathbf{x}}$.

D.5.2 Additional Thresholding

When utilizing SBL for line spectrum estimation—which involves the applications of arrival direction and arrival time estimation—it is known to overestimate the number of components [30–33]. A common solution to this problem is to introduce an additional thresholding operator that prunes away components for which little evidence exists in the data [32, 33]. For variational SBL, [13] suggested to introduce such a pruning as follows. In order to be a limit of the variational update sequence, each fixed point $\hat{\gamma}_i^* \in \mathbb{R}_{\geq 0} : f_i(\hat{\gamma}_i^*) - \hat{\gamma}_i^* = 0$ must be locally stable. Any point $\hat{\gamma}_i^*$ of the recurrence $f_i(\hat{\gamma}_i)$ is locally stable if, and only if [13]

$$\left| \frac{df_i(\hat{\gamma}_i)}{d\hat{\gamma}_i} \right|_{\hat{\gamma}_i=\hat{\gamma}_i^*} < 1. \quad (\text{D.35})$$

Reference [13] showed, that for $d_i = 1$ and using Jeffrey's prior as hyperprior PDF $p(\gamma_i)$, the condition (D.35) equals $\frac{q_{i,1}^2}{s_{i,1}} \geq 1$. Furthermore, [13] suggested to introduce a heuristic threshold $\frac{q_{i,1}^2}{s_{i,1}} > \chi' \geq 1$ to increase the sparsity of the estimator. They verified the effectiveness of this method through numerical simulations. For $d_i > 1$, the relation between condition (D.35) and the values of $q_{i,l}$ and $s_{i,l}$, $l \in \{1, 2, \dots, d_i\}$ is more involved. Nevertheless, we can apply the same intuition and heuristically introduce a threshold

$$\left| \frac{df_i(\hat{\gamma}_i)}{d\hat{\gamma}_i} \right|_{\hat{\gamma}_i=\hat{\gamma}_i^*} < \chi \leq 1 \quad (\text{D.36})$$

which each fixed point $\hat{\gamma}_i^*$ must fulfill in order to be considered as member of the sets of potential fixed points \mathcal{G}_i for the fast update rule. However, one disadvantage of this approach is that for $\chi < 1$ the fast update might not correspond to the result of an infinite sequence of variational updates anymore. Hence, any guarantee that the ELBO is increased in each step and that the algorithm converges is lost.

Let us point out, that for the scaled Jeffrey’s PDF with $c_i > 0$ increasingly more mass of the hyperprior PDF $p(\gamma_i)$ is placed on larger hyperparameters γ_i . Since the hyperparameters γ_i scale the prior precision matrix of each block, this can be interpreted as an additional force that drives the weight estimates towards $\hat{\mathbf{x}}_i = \mathbf{0}$. Hence, the parameter c_i can be used to tune the desired (i.e. the a priori expected) sparsity of the estimator in a similar fashion as the heuristic threshold χ . However, no modification of the update rules is required in this case. Therefore, the algorithm is still guaranteed to increase the ELBO in each step and, thus, guaranteed to converge. Note, that using a scaled Jeffrey’s prior with $c_i > 0$ results in a trivial maximum of the ELBO at $\gamma_i = \infty \forall i$, i.e. where $\hat{\mathbf{x}}_i = \mathbf{0} \forall i$. However, this trivial global maximum of the ELBO did not impact our numerical results, since the algorithm is only guaranteed to converge towards a local maximum. If the hyperparameters are initialized with $\hat{\gamma}_i = 0 \forall i$, then we observed that our algorithm converges to a non-trivial local maximum of the ELBO in our simulations. Analyzing the behavior of the algorithm under the existence of such a trivial maximum as well as designing priors which increase the sparsity without resulting in such a trivial global maximum is an interesting avenue for future research.

D.5.3 Simulation Study

In order to verify the theoretical analysis of Section D.5.1 and to investigate the heuristic approach of Section D.5.2, we conduct the following experiment that investigates the thresholding behaviour and the resulting sparsity of the proposed algorithm versus the choice of the hyperprior $p(\gamma_i)$. A noiseless system with a single block where $K = 1$, $d_1 = M = N$, and $\Phi = \mathbf{I}_N$ is considered, i.e., $\mathbf{y} = \mathbf{x} = \mathbf{x}_1$.

Let $\mathbf{1}_N$ denote a vector of ones with length N , noiseless measurements $\mathbf{y} = \alpha \mathbf{1}_N$ are generated repeatedly using different scales α . We evaluate the fast update rule for $\hat{\gamma}_i$ and the resulting weight estimate $\hat{\mathbf{x}}_i$ depending on the amplitude α of \mathbf{y} for $i = 1$ and fixed $\hat{\lambda} = 1$. Figure D.1a and D.1b show the RMS amplitude $\|\hat{\mathbf{x}}_i\|/\sqrt{d_i}$ of the weights $\hat{\mathbf{x}}_i$, where $\|\cdot\|$ denotes the L_2 -norm. “Hard” thresholding is given by $\hat{\mathbf{x}}_i = \mathbf{y}$ if $\|\mathbf{y}\|/\sqrt{d_i} > 1$ and $\hat{\mathbf{x}}_i = \mathbf{0}$ otherwise. It is shown by the dashed gray lines in Figure D.1.

Since for this setup $\Phi = \mathbf{I}_N$, a perfect estimator for the noiseless case would be $\hat{\mathbf{x}}_i = \mathbf{y}$ and a sparse estimate can be obtained by thresholding small values of $\|\hat{\mathbf{x}}_i\|$ for $i = 1$. Comparing the estimator obtained by different parameters of the generalized inverse Gaussian hyperprior PDF, it can be observed that the Gamma prior and the inverse Gamma prior do not lead to a sparse estimator since $\hat{\mathbf{x}}_i = \mathbf{0}$ if and only if $\mathbf{y} = \mathbf{0}$. This is in line with both the analysis presented in Section D.5.1 and the results from the literature [15, 19].⁸ When comparing the scaled Jeffrey’s prior and Jeffrey’s prior, we see that both estimators are indeed sparse since many values of \mathbf{y} lead to $\hat{\mathbf{x}}_i = \mathbf{0}$. Furthermore, we can see that Jeffrey’s prior without an

⁸Note that the non-blocked complex model presented in [19] can be viewed as a block-sparse model where each block consists of the real and imaginary part of each weight.

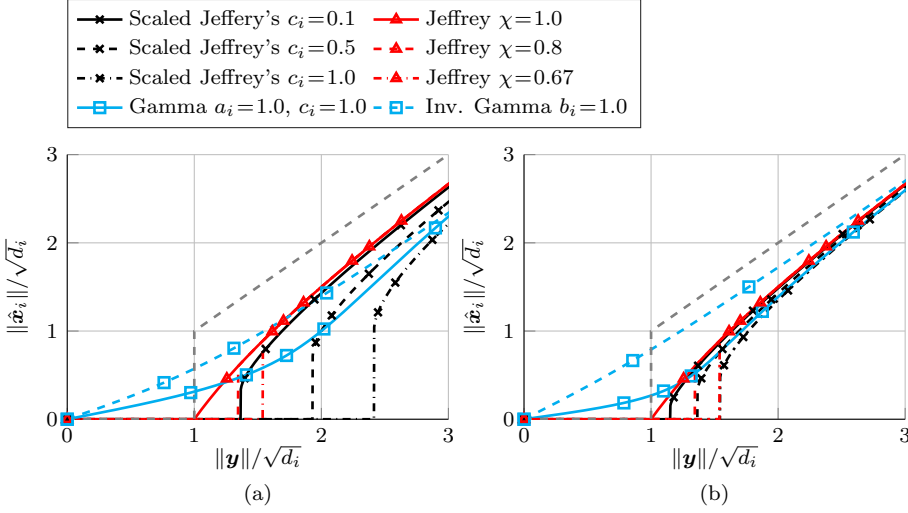


Figure D.1: Thresholding function for different priors and parameters for a system with a single block $K = 1$, $d_i = M = N$ and $\Phi = \mathbf{I}_N$. Using a Gamma PDF or inverse Gamma PDF with $b > 0$ as hyperprior does not result in a sparse estimate whereas using Jeffrey's prior or scaled Jeffrey's prior results in a sparse estimate. Block sizes $d_i = 2$ (a) and $d_i = 10$ (b).

increased heuristic threshold χ leads to a “soft” thresholding behavior that smoothly approaches $\hat{\mathbf{x}}_i = \mathbf{0}$ as \mathbf{y} gets smaller. Introducing a threshold of $\chi < 1$ or utilizing the scaled Jeffrey's prior is analogous to the “hard” thresholding approach, where the threshold is determined by either χ or c_i . An increase in the threshold results obviously into a sparser estimate. Furthermore, it is noteworthy to highlight the effect of difference different block sizes on the solution ($d_i = 2$ and $d_i = 10$ shown in Figure D.1a and D.1b, respectively). The cutoff value for Jeffrey's prior remains constant even with different thresholds, relative to the average amplitude of the components within the block. However, the likelihood of the average component amplitude being high decreases as the block size increases. This is due to the low likelihood of all amplitudes being high in the same realization. Consequently, a fixed threshold χ would result in varying false alarm rates depending on the block size. In contrast, using a scaled Jeffrey's prior leads to a reduction of the cutoff region with an increase in block size, compensating to some extent for the reduction of the false alarm rate.

Finally, from Figure D.1 we can see that with c_i and χ tuned to achieve a similar thresholding (e.g. $c_i = 1$ and $\chi = 0.67$ for a block size of $d_i = 10$), the estimation performance of both estimators is rather similar. Thus, we conclude that the scaled Jeffrey's prior is to be preferred over Jeffrey's prior with an additional threshold in order to retain the convergence guarantees of the variational algorithm.

D.5.4 Equivalence to Fast Marginal Likelihood Maximization

In the following, we prove that the proposed fast variational solution to block-SBL using Jeffrey's prior is equivalent to the fast marginal likelihood maximization

scheme of [12]. In [12], the hyperparameters γ are estimated by maximizing the log marginal likelihood

$$\mathcal{L}_{\text{Ma}}(\gamma) = 2 \ln \int p(\mathbf{y}|\mathbf{x}, \lambda) p(\mathbf{x}|\gamma) d\mathbf{x} \quad (\text{D.37})$$

with respect to a single γ_i at a time in an iterative fashion. We start to prove the equivalence between [12] and our solution by expressing the marginal likelihood (D.37) as

$$\mathcal{L}_{\text{Ma}}(\gamma) = \ln |\Sigma| + \hat{\lambda}^2 \mathbf{y}^H \Phi \Sigma \Phi^H \mathbf{y} + \sum_{i=1}^K d_i \ln \gamma_i + \text{const.} \quad (\text{D.38})$$

after integrating out the weights \mathbf{x} . As derived in Appendix D.9.2, (D.38) can be expressed as a function of a single prior γ_i using the Woodbury matrix identity and block matrix determinant lemma as

$$\begin{aligned} \mathcal{L}_{\text{Ma}}(\gamma) &= \mathcal{L}_{\sim i}(\gamma_{\sim i}) + \mathcal{L}_i(\gamma_i) \\ &= \mathcal{L}_{\sim i}(\gamma_{\sim i}) + \sum_{l=1}^{d_i} \ln \left(\frac{\gamma_i s_{i,l}}{1 + \gamma_i s_{i,l}} \right) - \frac{\gamma_i |q_{i,l}|^2}{(1 + \gamma_i s_{i,l})^2} \end{aligned} \quad (\text{D.39})$$

where $\gamma_{\sim i}$ denotes the vector γ with the element γ_i removed and $\mathcal{L}_{\sim i}(\gamma_{\sim i})$ is some function which does not depend on γ_i . The marginal likelihood is maximized by finding the maximum of (D.39) as $\frac{d}{d\gamma_i} \mathcal{L}_i(\gamma_i) = 0$. The derivative of (D.39) is

$$\frac{d}{d\gamma_i} \mathcal{L}_i(\gamma_i) = \sum_{l=1}^{d_i} \frac{1}{\gamma_i} \frac{1}{1 + \gamma_i s_{i,l}} - \frac{|q_{i,l}|^2}{(1 + \gamma_i s_{i,l})^2} \quad (\text{D.40})$$

Thus, the condition $\frac{d}{d\gamma_i} \mathcal{L}_i(\gamma_i) = 0$ can be expressed as

$$\begin{aligned} 0 &= \gamma_i \sum_{l=1}^{d_i} \left((\gamma_i s_{i,l}^2 + |q_{i,l}|^2 + s_{i,l}) \prod_{j=1, j \neq l}^{d_i} (1 + \gamma_i s_{i,j}) \right) \\ &\quad - d_i \prod_{l=1}^{d_i} (1 + \gamma_i s_{i,l})^2 \\ &= \gamma_i B_i(\gamma_i) - d_i A_i(\gamma_i). \end{aligned} \quad (\text{D.41})$$

applying the definitions (D.33) and (D.34). Comparing (D.41) with the polynomials $G_i(\hat{\gamma}_i)$ in the fifth column of Table D.1, we recognize that the extrema of the marginalized likelihood with respect to a single γ_i correspond to the fixed points of the recurrent relation f_i in the fast variational solution in case Jeffrey's prior is used. Hence, showing that the two approaches are equivalent.

D.6 Numerical Evaluation

To investigate the performance of the algorithm, we generate a dictionary matrix with $M = 2N$ columns and assume $N = 200$ measurements are obtained unless otherwise stated. The elements of the dictionary matrix are drawn from independent

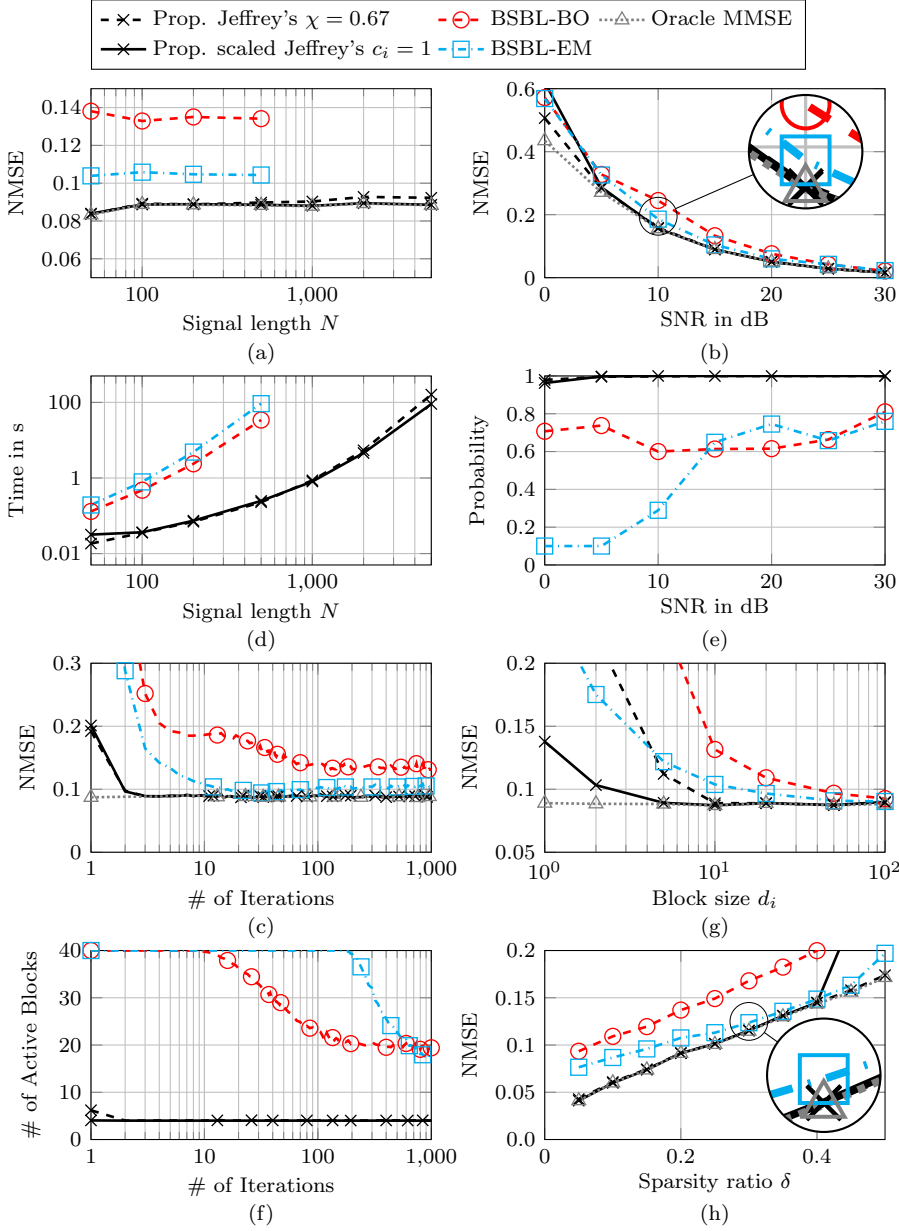


Figure D.2: Performance of the proposed algorithm using a dictionary matrix with Gaussian distributed entries averaged over 100 realizations. NMSE as function of the signal length (a), SNR (b) and the number of iterations (c). The runtime of each algorithm as a function of the signal length is shown in (d). Classification probability of correctly classifying each block as zero/nonzero as a function of the SNR (e) and the number of active blocks ($\hat{\mathbf{x}}_i \neq \mathbf{0}$) as function of the number of iterations (f). Further evaluation of the NMSE over block size (g) and sparsity ratio (h). Unless otherwise stated, the following parameters were used in the simulations. $N = 200$, $M = 2N$, SNR = 15 dB, $\delta = 0.2$ and $d_i = 10 \forall i$.

zero-mean normal distributions with unit variance and each column ϕ_k of Φ is normalized such that $\|\phi_k\| = 1$. The corresponding weight vector \mathbf{x} is partitioned into blocks with an equal block size of $d_i = 10 \forall i$. We selected a number of blocks at randomly chosen locations such that the desired sparsity ratio $\delta = \frac{\|\mathbf{x}\|_0}{N} = 0.2$ was achieved. For these nonzero blocks we draw \mathbf{x}_i from a Gaussian distribution with zero mean and unit variance while for the remaining blocks $\mathbf{x}_i = 0$. The location of the nonzero blocks is unknown to the algorithm. The noise precision λ is chosen, such that the signal-to-noise ratio (SNR) defined as $\text{SNR} = \frac{\lambda \|\Phi \mathbf{x}\|^2}{N}$ equals 15 dB. As performance metric we use the normalized mean squared error (NMSE) defined as $\text{NMSE} = \frac{\|\mathbf{x} - \hat{\mathbf{x}}\|^2}{\|\mathbf{x}\|^2}$, averaged over 100 simulation runs. We use Jeffrey's prior for the noise precision $p(\lambda)$, which is obtained by $\epsilon = \eta = 0$ in (D.7). Two variants of the proposed method are evaluated. Once using a Jeffrey's prior, obtained by $a_i = b_i = c_i = 0$ in (D.15), as hyperprior $p(\gamma_i)$ and a threshold of $\chi = 0.67$, and once using the scaled Jeffrey's prior with $c_i = 1$ as hyperprior $p(\gamma_i)$. Since we are interested in sparse estimators, we did not consider the non-sparse variants of the hyperprior. As comparison methods we use two variants of the EM-based block-SBL algorithm from [9].⁹ The comparison algorithms apply the EM algorithm directly (BSBL-EM) or use a bounded-optimization approach to increase the convergence speed of the EM algorithm (BSBL-BO). Note, that the BSBL-EM algorithm is equivalent to the variational BSBL algorithm [15] using Jeffrey's prior as hyperprior. For all algorithms we assumed that $\mathbf{B}_i = \mathbf{I}$ is a known and fixed parameter. We also include an "oracle" minimum mean squared error (MMSE) estimator, which is given the true locations of the nonzero blocks and calculates the weights of the nonzero blocks using the pseudoinverse of the corresponding columns of the dictionary as reference.

Figure D.2a-D.2c depicts the NMSE of the algorithms as a function of the signal length N , the SNR and number of iterations. We define the number of iterations as the number of times the main loop of the algorithm is repeated, i.e. the amount of times all variables $\hat{\mathbf{x}}$, $\hat{\Sigma}$, $\hat{\lambda}$ and $\hat{\gamma}_1 \dots \hat{\gamma}_K$ are updated. Figure D.2d shows the runtime of the algorithms again as function of the signal length N . For $N = 500$, the proposed method is faster by approximately 2 orders of magnitude compared to the two BSBL variants [9] while achieving an NMSE virtually identical to that of the oracle MMSE estimator. As shown in Figures D.2b and D.2c, an NMSE almost identical to that of the oracle estimator is achieved already after 3 iterations and over a wide range of SNRs. Figure D.2e shows the empirical probability of correctly estimating any $\hat{\gamma}_i$ as converged or diverged (i.e. the empirical probability of classifying each block correctly as zero or nonzero block). While the proposed method classifies the blocks correctly as zero or nonzero almost all the time, both variants of the BSBL algorithm significantly overestimated the number of nonzero blocks.

The difference in the runtime of the algorithms can be explained by Figures D.2c and D.2f. While our proposed algorithm converges to an NMSE estimate virtually identical that of the oracle estimator after 3 iterations, the BSBL-BO and BSBL-EM algorithm require more iterations to converge. Furthermore, as detailed in Figure D.2f, the proposed algorithm estimates the number of active

⁹The code for the BSBL algorithm was obtained from <http://dsp.ucsd.edu/~zhilin/BSBL.html>

blocks correctly from the first iteration on. Thus, any operations that require the covariance of the weights $\hat{\Sigma}$ can be performed with a small matrix since the entries of $\hat{\Sigma}$ corresponding to deactivated blocks are exactly zero. It can be seen in Figure D.2f, that the BSBL-BO algorithm keeps all 40 blocks active for the first 10 iterations and then slowly starts to deactivate blocks. The BSBL-EM algorithm keeps all 40 blocks active for the first 100 iterations before it starts to deactivate them. Thus, in addition to requiring more iterations until convergence, each iteration also involves operations with a larger matrix $\hat{\Sigma}$. Therefore, each iteration of both BSBL variants is computationally more expensive compared to the proposed algorithm.

Additionally, we evaluate how varying the block size d_i and the sparsity ratio δ affects the algorithm performance in terms of NMSE. To evaluate the performance of the algorithm as function of the block size d_i , we use $N = 500$ while we use $N = 200$ for the evaluation over the sparsity ratio δ . Figures D.2g and D.2h show the NMSE obtained by these numerical experiments. Again, the proposed method outperforms both BSBL variants over most simulated sparsity ratios and SNRs. An evaluation of the runtime and zero/nonzero classification probability was evaluated as well. However we found these results to be similar to the ones in Figure D.2d and D.2e. Thus, they are omitted for the sake of brevity. In case the block size is varied, the proposed method utilizing the scaled Jeffrey's prior with $c_i = 1$ is again superior to both BSBL variants as well as to the proposed method if Jeffrey's prior is used in combination with a threshold of $\chi = 0.67$. The difference between the two variants of the proposed method stems from the dependence of the probability of missed detection and false alarm on the block size, which is more severe for Jeffrey's prior compared to the scaled Jeffrey's prior as discussed in Section D.5.3.

D.7 Application Example: DOA Estimation

D.7.1 System Model

Consider an unknown number of L narrowband sources with complex signal amplitudes $s_l[t] \in \mathbb{C}$, $l = 1, \dots, L$ located at fixed DOAs $\theta_l \in \Theta = [-90^\circ, 90^\circ)$ in the far field of an array consisting of N sensors, such as a microphone or antenna array. Multiple measurements $\mathbf{y}_t \in \mathbb{C}^N$ at different times $t = 1, 2, \dots, d$ are obtained. Each measurement is modeled as

$$\mathbf{y}_t = \sum_{l=1}^L \boldsymbol{\psi}(\theta_l) s_l[t] + \mathbf{v}_t \quad (\text{D.42})$$

where $\boldsymbol{\psi}(\theta_l)$ is the steering vector of the array and \mathbf{v}_t is additive noise. We assume that the noise \mathbf{v}_t follows a Gaussian PDF $\mathcal{N}(\mathbf{v}_t; \mathbf{0}, \lambda^{-1} \mathbf{I})$ and is independent across different times t . For a linear array, the steering vector is

$$\boldsymbol{\psi}(\theta_l) = \frac{1}{\sqrt{N}} \left[1 \ e^{-j2\pi \frac{p_2}{\mu} \sin \theta_l} \ \dots \ e^{-j2\pi \frac{p_N}{\mu} \sin \theta_l} \right]^T \quad (\text{D.43})$$

where μ is the wavelength of the signal and p_n is the distance from sensor 1 to sensor n .

An alternative model can be obtained by introducing a search grid of K potential source DOAs $\bar{\boldsymbol{\theta}} = [\bar{\theta}_1 \ \bar{\theta}_2 \ \cdots \ \bar{\theta}_K]^T$ and corresponding amplitude vectors $\mathbf{x}^{(t)}$. With this grid we construct a dictionary matrix $\boldsymbol{\Psi} = [\boldsymbol{\psi}(\bar{\theta}_1) \ \boldsymbol{\psi}(\bar{\theta}_2) \ \cdots \ \boldsymbol{\psi}(\bar{\theta}_K)]$, such that we approximate (D.42) as

$$\mathbf{y}_t = \boldsymbol{\Psi} \mathbf{x}^{(t)} + \mathbf{v}_t. \quad (\text{D.44})$$

Instead of directly estimating L and the DOAs θ_1 through θ_L we obtain a sparse estimate of $\mathbf{x}^{(t)}$ and use the number of nonzero elements of $\mathbf{x}^{(t)}$ as estimate \hat{L} of L . From the columns of $\boldsymbol{\Psi}$ corresponding to the active entries in $\mathbf{x}^{(t)}$, we can extract DOAs estimates $\hat{\theta}_l$, $l = 1, \dots, \hat{L}$.

The locations of the sources are assumed stationary. Therefore, the sparsity profile of $\mathbf{x}^{(t)}$ at different times t stays constant. Let $\mathbf{Y} = [\mathbf{y}_1 \ \mathbf{y}_2 \ \cdots \ \mathbf{y}_d]$, $\mathbf{X} = [\mathbf{x}^{(1)} \ \mathbf{x}^{(2)} \ \cdots \ \mathbf{x}^{(d)}]$ and $\mathbf{V} = [\mathbf{v}_1 \ \mathbf{v}_2 \ \cdots \ \mathbf{v}_d]$, we arrive at the row-sparse MMV signal model

$$\mathbf{Y} = \boldsymbol{\Psi} \mathbf{X} + \mathbf{V}. \quad (\text{D.45})$$

As discussed in Section D.2.1, (D.45) can be rearranged into a block-sparse model. Thus, the row-sparse matrix \mathbf{X} can be estimated using the presented algorithm.

Note, that in this model the weights in each row of \mathbf{X} , (i.e. each block \mathbf{x}_i of \mathbf{x}) are either zero or correspond to the amplitudes of a source over time $\mathbf{x}_i = [s_i[1] \ s_i[2] \ \cdots \ s_i[d]]^T$. Thus, any prior knowledge about these source amplitudes can be incorporated in the matrix \mathbf{B}_i .

D.7.2 DOA Estimation Results

Since the SBL algorithms from [9] are not applicable for a complex valued signal model, we compare our algorithm against the multi-snapshot SBL-based DOA estimation algorithm in [20] termed DOA-SBL.¹⁰ Since in this case the actual weights \mathbf{X} are of secondary interest, we resort to the optimal subpattern assignment (OSPA) metric [34] of the estimated DOAs $\hat{\theta}_l$, $l \in \{1, \dots, \hat{L}\}$ to evaluate the performance of the algorithms. The OSPA is a metric which considers both the actual estimation error as well as the error in estimating the model order (i.e. the number of missed detection and false alarms). As parameters for the OSPA, we set the order to 1 and use a cutoff-distance of 5° . Since the DOA-SBL algorithm does not directly estimating a sparse weight vector, we evaluate two versions of this algorithm. Let $\gamma_{i,\text{DOA-SBL}}$ and $\sigma_{\text{DOA-SBL}}^2$ denote the prior variances and noise variance estimated by [20]. We define the component SNR of the i -th block as $\text{SNR}_{i,\text{DOA-SBL}} = \frac{\gamma_{i,\text{DOA-SBL}}}{\sigma_{\text{DOA-SBL}}^2}$. We consider all DOAs $\hat{\theta}_i$ which have an associated component SNR above a certain threshold as first variant of the DOA-SBL algorithm.¹¹ Based on preliminary simulations, we adapted the threshold to maximize the performance of the DOA-SBL algorithm in our simulated scenarios. As an additional comparison, we also consider an oracle variant of the algorithm

¹⁰The code for the DOA-SBL algorithm was obtained from <https://github.com/gerstoft/SBL>.

¹¹The authors of the DOA-SBL algorithm used a similar thresholding based on the maximum estimated prior variance $\gamma_{i,\text{DOA-SBL}}$ in a related paper [35]. However, additional simulations which are omitted for brevity show that our thresholding based on the component SNR achieves a smaller OSPA than the method used in [35].




Figure D.3: Runtime comparison (a) and OSPA (b) of the proposed method with the DOA-SBL algorithm [20] as function of the number of array elements at an $\text{SNR}_A = 10$ dB given a single snapshot $d = 1$.

which is given the true number of components L and estimates $\hat{\theta}_l$ as the DOAs corresponding to the L blocks with largest component SNR. To account for the increasing processing gain at increasing arrays sizes, we define the array SNR as for snapshot t as $\text{SNR}_A = \lambda \|\Psi \mathbf{x}^{(t)}\|^2$.

Single-Snapshot Estimation Performance

For the first experiment, we analyse the runtime and estimation accuracy as a function of the number of array elements N . We generate a dictionary with $M = 2N$ entries spaced such that $\sin(\bar{\theta}_m)$ forms a regular grid. We simulate 3 sources located at the grid points closest to $\{-2^\circ, 3^\circ, 50^\circ\}$ with amplitudes drawn randomly from a zero-mean complex Gaussian distribution with unit variance and consider only a single snapshot $d = 1$. For the DOA-SBL algorithm, we calculate the OSPA based on all blocks with component SNR ≥ 10 dB. To illustrate the effect of the parameter c_i , we use two variants of our algorithm. Both use the scaled Jeffrey's priors and are parameterized with $c_i = 1$ and $c_i = 2$, respectively. Figures D.3a and D.3b show the runtime and OSPA, respectively, for both algorithms at an $\text{SNR}_A = 20$ dB. The smallest OSPA is achieved by the oracle DOA-SBL algorithm. This is not surprising, since this algorithm is given the true number of sources in advance. However, knowing the true number of sources is not a realistic assumption for most




Figure D.4: OSPA of the proposed method and the DOA-SBL algorithm [20] as function of the SNR_A given $d = 10$ snapshots with different source correlations β .

practical applications. For array sizes of 100 elements and larger, the proposed algorithm parameterized with $c_i = 2$ is faster than the DOA-SBL algorithm while achieving the same OSPA. For array sizes with less than 100 elements, the OSPA for our algorithm parameterized with $c_i = 2$ increases rapidly. This is due to some components not being detected in those cases. The parameter c_i acts similar to a threshold. Hence, we can increase the detection probability by using a smaller value for c_i . Consequently, this also leads to an increased false alarm rate. The increased false alarm rate is mostly noticeable at large N , due to the larger number of possible DOAs in the grid for large N . Hence, by using a smaller c_i we can reduce the OSPA at smaller array sizes at the cost of increasing the OSPA at larger array sizes, as shown in Figure D.3b. We refer the reader to [32] for a detailed discussion on the relation between the array size and number of false alarms for a similar SBL-based approach.

Multi-Snapshot Estimation Performance

Next, we simulate a system with an array consisting of $N = 100$ antennas where $d = 10$ snapshots are obtained. We simulate three sources at the same DOAs as in the previous experiment. To investigate the effect of intra-block correlations, the source amplitudes $s_k[t]$ are generated by a first-order autoregressive model, i.e. $s_k[t] = \xi + \beta s_k[t - 1]$ where $\xi \sim \mathcal{CN}(\xi; 0, 1)$ is a circularly-symmetric zero-mean complex Gaussian random variable with unit variance. The coefficient $\beta \in \mathbb{C} : |\beta| < 1$ can be used to set the correlation between the snapshots over time. The

covariance matrix Σ_s of the samples $s_k[t]$ is given as a Toeplitz matrix

$$\Sigma_s = \begin{bmatrix} 1 & \beta & \dots & \beta^{d-1} \\ \beta & 1 & \dots & \beta^{d-2} \\ \vdots & \vdots & \ddots & \vdots \\ \beta^{d-1} & \beta^{d-2} & \dots & 1 \end{bmatrix}. \quad (\text{D.46})$$

We evaluate three different cases. (i) no correlation $\beta = 0$, (ii) medium correlation $\beta = 0.5$ and (iii) strong correlation $\beta = 0.95$. For the proposed algorithm we set $\mathbf{B}_i = \Sigma_s^{-1} \forall i$ to exploit this information. The correlation between the source amplitudes provides additional statistical information which can be used to separate true sources from the additive white noise. Therefore, the probability of false alarms gets smaller with increasing correlation β . To achieve an (approximately) constant false alarm rate for our algorithm, we parameterize the scaled Jeffrey's prior with $c_i = 2$ for the case of no correlation (i) and medium correlation (ii), and use $c_i = 1.5$ for high correlation (iii), based on preliminary simulations. For the DOA-SBL algorithm, we calculate the OSPA based on all blocks with a component $\text{SNR} \geq 2\text{ dB}$, based on the same preliminary simulations. The performance of the DOA-SBL algorithm is approximately the same in all three cases, since it can not exploit the additional information resulting from the correlation of the source amplitudes. Thus, we plot the performance of the DOA-SBL only for case (i). Figure D.4 depicts the OSPA of both algorithms as a function of the SNR_A . For $\text{SNR}_A < 0\text{ dB}$ and $\text{SNR}_A > 10\text{ dB}$, the estimation either fails due to the high noise level or is trivial. Thus the performance of both algorithms is approximately the same in these regions. In the transition region $0\text{ dB} \leq \text{SNR}_A \leq 10\text{ dB}$, the proposed algorithm achieves a smaller OSPA than the DOA-SBL algorithm for all three cases. With increasing correlation, the performance of the proposed algorithm improves. For the case of high correlation (iii), the performance of the proposed algorithm is practically the same as the performance of the oracle DOA-SBL algorithm which is given the true number of sources L . We refer the reader to [8] for a more in-depth investigation and discussion of the effects of the intra-block correlation \mathbf{B}_i as well as for suggestions on how to estimate this matrix efficiently without introducing too many additional parameters.

Note, that in any practical example the DOAs will not align exactly with the search grid of potential DOAs. This mismatch introduces errors and complicates the estimation process. This can be counteracted e.g. by using a variational-EM approach similar to [14, 21] to optimizing the ELBO over the estimated source locations $\hat{\theta}_l$ in addition to the variational parameters. Furthermore, [30] directly integrates the estimation of the parameters $\hat{\theta}_l$ into the variational framework in order to obtain (approximate) posterior distributions $q(\theta_l)$. However, we consider these off-grid approaches outside the scope of this work.

D.8 Conclusion

We derived a fast update rule for the variational Bayesian approach to block-SBL. Iterating the update equation for the approximating PDF of the weights and hyperparameters is expressed as a first-order recurrence from the previous parameter set to the next. We showed how the fixed points of the recurrent relation

can be obtained as the roots of a polynomial. Furthermore, by proving that the recurrent relation is a strictly increasing rational function, we are able to determine if the sequence converges and, if it does converge, we can determine its limit. Hence, we can check for convergence/divergence and update each hyperparameter to its asymptotic value in a single step.

Using numerical simulations, we show that the proposed fast update rule improves the run time of the variational block-sparse SBL algorithm [9] by two orders of magnitude. Moreover, the proposed algorithm is shown to achieve a smaller NMSE over a wide range of signal lengths, SNRs, block sizes and sparsity ratios than the comparison algorithms. Additionally, the proposed algorithm is also applied to DOA estimation. Our results demonstrate that (i) our algorithm is able to outperform similar SBL-based DOA estimation algorithms in terms of computation time while achieving a similar performance for single-snapshot DOA estimation and (ii) achieves a smaller OSPA than the comparison DOA estimation algorithm in a low-SNR multiple-snapshot DOA estimation scenario where the source amplitudes are correlated over the snapshots.

A promising direction for future research is the extension of the presented method to include the estimation of the size of each block, or, to a parametrized continuous (infinite) dictionary matrix similar to our recent work [21]. Another interesting direction for future research is the combination of the presented method with belief propagation algorithms for joint channel estimation and decoding [25] or sequential tracking of time variant channels [36] in multi-input multiple-output communication systems, which are often modeled to be block-sparse [23, 24].

D.9 Appendix

D.9.1 Proof that h_i is Strictly Decreasing

To obtain a more convenient expression for $h_i(\hat{\gamma}_i) = B_i(\hat{\gamma}_i)/A_i(\hat{\gamma}_i)$, we express the polynomial $B_i(\hat{\gamma}_i)$ given in (D.34) as the product of $A_i(\hat{\gamma}_i)$ defined in (D.33) and some expression

$$\begin{aligned} B_i(\hat{\gamma}_i) &= \sum_{l=1}^{d_i} (\hat{\gamma}_i s_{i,l}^2 + |q_{i,l}|^2 + s_{i,l}) \prod_{j=1, j \neq i}^{d_i} (1 + \hat{\gamma}_i s_{i,j})^2 \\ &= \sum_{l=1}^{d_i} \frac{\hat{\gamma}_i s_{i,l}^2 + |q_{i,l}|^2 + s_{i,l}}{(1 + \hat{\gamma}_i s_{i,l})^2} \prod_{j=1}^{d_i} (1 + \hat{\gamma}_i s_{i,j})^2 \\ &= A_i(\hat{\gamma}_i) \cdot \sum_{l=1}^{d_i} \frac{\hat{\gamma}_i s_{i,l}^2 + |q_{i,l}|^2 + s_{i,l}}{(1 + \hat{\gamma}_i s_{i,l})^2}. \end{aligned} \quad (\text{D.47})$$

Inserting (D.47) into (D.21), we rewrite $h_i(\hat{\gamma}_i)$ as

$$h_i(\hat{\gamma}_i) = \frac{B_i(\hat{\gamma}_i)}{A_i(\hat{\gamma}_i)} = \sum_{l=1}^{d_i} \frac{\hat{\gamma}_i s_{i,l}^2 + |q_{i,l}|^2 + s_{i,l}}{(1 + \hat{\gamma}_i s_{i,l})^2}. \quad (\text{D.48})$$

Taking the derivative of $h_i(\hat{\gamma}_i)$ with respect to $\hat{\gamma}_i$ results in

$$\begin{aligned} \frac{d}{d\hat{\gamma}_i} h_i(\hat{\gamma}_i) &= \sum_{l=1}^{d_i} \frac{d}{d\hat{\gamma}_i} \frac{\hat{\gamma}_i s_{i,l}^2 + |q_{i,l}|^2 + s_{i,l}}{(1 + \hat{\gamma}_i s_{i,l})^2} \\ &= - \sum_{l=1}^{d_i} \frac{\hat{\gamma}_i s_{i,l}^3 + 2s_{i,l}|q_{i,l}|^2 + s_{i,l}^2}{(1 + \hat{\gamma}_i s_{i,l})^3}. \end{aligned} \quad (\text{D.49})$$

The variables $s_{i,l}$ are the eigenvalues of a positive definite matrix and, thus, $s_{i,l} > 0$, resulting in $\frac{dh_i(\hat{\gamma}_i)}{d\hat{\gamma}_i} < 0 \forall \hat{\gamma}_i \in \mathbb{R}_{\geq 0}$. Thus, proving that h_i is a strictly decreasing function on the domain $\mathbb{R}_{\geq 0}$.

D.9.2 Detailed Derivations for the Proof of Equivalence

Let Φ_i denote all columns of the dictionary matrix which correspond to the i -th block and $\Phi_{\sim i}$ all remaining columns of Φ , such that $\Phi = [\Phi_{\sim i} \ \Phi_i]$. Similarly, let $B_{\sim i}$ be a diagonal matrix with the blocks $\gamma_k B_k \forall k \in \{1, 2, \dots, K\} \setminus \{i\}$ on its main diagonal, such that we can express Σ^{-1} as a block matrix¹²

$$\Sigma^{-1} = \begin{bmatrix} \lambda \Phi_{\sim i}^H \Phi_{\sim i} + B_{\sim i} & \lambda \Phi_{\sim i}^H \Phi_i \\ \lambda \Phi_i^H \Phi_{\sim i} & \lambda \Phi_i^H \Phi_i + \gamma_i I \end{bmatrix}. \quad (\text{D.50})$$

Denoting $\Sigma_{\sim i} = (\lambda \Phi_{\sim i}^H \Phi_{\sim i} + B_{\sim i})^{-1}$ and $\Sigma_i = (\lambda \Phi_i^H \Phi_i - \lambda^2 \Phi_i^H \Phi_{\sim i} \Sigma_{\sim i} \Phi_{\sim i}^H \Phi_i)^{-1}$, we use the block matrix determinant lemma to express

$$\begin{aligned} \ln |\Sigma| &= -\ln |\Sigma^{-1}| = \ln |\Sigma_{\sim i}| - \ln |\Sigma_i^{-1} + \gamma_i I| \\ &= \ln |\Sigma_{\sim i}| - \sum_{l=1}^{d_i} \ln \left(\gamma_i + \frac{1}{s_{i,l}} \right). \end{aligned} \quad (\text{D.51})$$

Using the Woodbury matrix identity, we can also make the dependence on γ_i explicit in the second term in (D.38)

$$\lambda^2 \mathbf{y}^H \Phi \Sigma \Phi^H \mathbf{y} = \lambda^2 \mathbf{y}^H \Phi (\bar{\Sigma} - \bar{\Sigma} E_i (\gamma_i^{-1} I + \Sigma_i)^{-1} E_i^T \bar{\Sigma}) \Phi^H \mathbf{y} \quad (\text{D.52})$$

Using the identity (D.28), expression (D.52) can be further simplified to

$$\begin{aligned} \lambda^2 \mathbf{y}^H \Phi \Sigma \Phi^H \mathbf{y} &= \text{const.} - \mathbf{q}_i^H (\gamma_i^{-1} I + \mathbf{S}_i)^{-1} \mathbf{q}_i \\ &= \text{const.} - \sum_{l=1}^{d_i} \frac{\gamma_i |q_{i,l}|^2}{1 + \gamma_i s_{i,l}}. \end{aligned} \quad (\text{D.53})$$

Inserting (D.51) and (D.53) into (D.38) we arrive at (D.39).

Given the derivatives $\frac{d}{d\gamma_i} \ln \frac{\gamma_i s_{i,l}}{1 + \gamma_i s_{i,l}} = \frac{1}{\gamma_i} \frac{1}{1 + \gamma_i s_{i,l}}$ and $\frac{d}{d\gamma_i} \frac{\gamma_i |q_{i,l}|^2}{1 + \gamma_i s_{i,l}} = \frac{|q_{i,l}|^2}{(1 + \gamma_i s_{i,l})^2}$, (D.40) is readily obtained. To arrive at (D.41), we start with $\frac{d}{d\gamma_i} \mathcal{L}_i(\gamma_i) = 0$ and multiply both sides of the equation with $\gamma_i \prod_{l=1}^{d_i} (1 + \gamma_i s_{i,l})^2$ to obtain

$$0 = \sum_{l=1}^{d_i} (1 + \gamma_i s_{i,l} - \gamma_i |q_{i,l}|^2) \prod_{j=1, j \neq l}^{d_i} (1 + \gamma_i s_{i,j})^2. \quad (\text{D.54})$$

¹²We use $B_i = I$ to simplify the notation, but the results can be easily verified to applied to the case of a general B_i as well.

Finally, we add $0 = \gamma_i s_{i,l} + \gamma_i^2 s_{i,l}^2 - \gamma_i s_{i,l} - \gamma_i^2 s_{i,l}^2$ into each term in the sum in (D.54) to arrive at (D.41) after a few algebraic manipulations.

Bibliography

- [1] M. F. Duarte and Y. C. Eldar, “Structured compressed sensing: From theory to applications,” *IEEE Trans. Signal Process.*, Sep. 2011, doi: 10.1109/TSP.2011.2161982.
- [2] O. Barndorff-Nielsen, J. Kent, and M. Sørensen, “Normal variance-mean mixtures and z distributions,” *Int. Statist. Rev. / Revue Internationale de Statistique*, vol. 50, no. 2, pp. 145–159, Aug. 1982.
- [3] M. E. Tipping, “Sparse Bayesian learning and the relevance vector machine,” *J. Mach. Learn. Res.*, vol. 1, pp. 211–244, Jun. 2001.
- [4] A. Faul and M. Tipping, “Analysis of sparse Bayesian learning,” in *Advances Neural Inf. Process. Syst.*, vol. 14, Vancouver, Canada, Dec. 3–8, 2001, pp. 383–389.
- [5] M. E. Tipping and A. C. Faul, “Fast marginal likelihood maximisation for sparse Bayesian models,” in *Proc. 9th Int. Workshop Artif. Intell. and Statist.*, vol. R4, Key West, FL, USA, Jan. 03–06, 2003, pp. 276–283.
- [6] D. P. Wipf and B. D. Rao, “Sparse Bayesian learning for basis selection,” *IEEE Trans. Signal Process.*, Aug. 2004, doi: 10.1109/TSP.2004.831016.
- [7] —, “An empirical Bayesian strategy for solving the simultaneous sparse approximation problem,” *IEEE Trans. Signal Process.*, Jun. 2007, doi: 10.1109/TSP.2007.894265.
- [8] Z. Zhang and B. D. Rao, “Sparse signal recovery with temporally correlated source vectors using sparse Bayesian learning,” *IEEE J. Sel. Topics Signal Process.*, 2011, doi: 10.1109/JSTSP.2011.2159773.
- [9] —, “Extension of SBL algorithms for the recovery of block sparse signals with intra-block correlation,” *IEEE Trans. Signal Process.*, Apr. 2013, doi: 10.1109/TSP.2013.2241055.
- [10] J. Fang, Y. Shen, H. Li, and P. Wang, “Pattern-coupled sparse Bayesian learning for recovery of block-sparse signals,” *IEEE Trans. Signal Process.*, Jan. 2015, doi: 10.1109/TSP.2014.2375133.
- [11] M. Luessi, S. D. Babacan, R. Molina, and A. K. Katsaggelos, “Bayesian simultaneous sparse approximation with smooth signals,” *IEEE Trans. Signal Process.*, Nov. 2013, doi: 10.1109/TSP.2013.2280441.
- [12] Z. Ma, W. Dai, Y. Liu, and X. Wang, “Group sparse Bayesian learning via exact and fast marginal likelihood maximization,” *IEEE Trans. Signal Process.*, May 2017, doi: 10.1109/TSP.2017.2675867.

- [13] D. Shutin, T. Buchgraber, S. R. Kulkarni, and H. V. Poor, "Fast variational sparse Bayesian learning with automatic relevance determination for superimposed signals," *IEEE Trans. Signal Process.*, Dec. 2011, doi: 10.1109/TSP.2011.2168217.
- [14] D. Shutin, W. Wand, and T. Jost, "Incremental sparse Bayesian learning for parameter estimation of superimposed signals," in *10th Int. Conf. Sampling Theory and Appl.*, Bremen, Germany, Jul. 1–5, 2013, pp. 513–516.
- [15] S. D. Babacan, S. Nakajima, and M. N. Do, "Bayesian group-sparse modeling and variational inference," *IEEE Trans. Signal Process.*, Jun. 2014, doi: 10.1109/TSP.2014.2319775.
- [16] S. Sharma, S. Chaudhury, and Jayadeva, "Variational Bayes block sparse modeling with correlated entries," in *24th Int. Conf. Pattern Recognit.*, Aug. 20–24, 2018, doi: 10.1109/ICPR.2018.8545302.
- [17] D. G. Tzikas, A. C. Likas, and N. P. Galatsanos, "The variational approximation for Bayesian inference," *IEEE Signal Process. Mag.*, Nov. 2008, doi: 10.1109/MSP.2008.929620.
- [18] C. M. Bishop, *Pattern Recognition and Machine Learning*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2006.
- [19] N. L. Pedersen, C. Navarro Manchón, M.-A. Badiu, D. Shutin, and B. H. Fleury, "Sparse estimation using Bayesian hierarchical prior modeling for real and complex linear models," *Signal Process.*, Oct. 2015, doi: <https://doi.org/10.1016/j.sigpro.2015.03.013>.
- [20] P. Gerstoft, C. F. Mecklenbräuker, A. Xenaki, and S. Nannuru, "Multisnapshot sparse Bayesian learning for DOA," *IEEE Signal Process. Lett.*, Oct. 2016, doi: 10.1109/LSP.2016.2598550.
- [21] J. Möderl, F. Pernkopf, K. Witrisal, and E. Leitinger, "Variational inference of structured line spectra exploiting group-sparsity," *ArXiv e-prints*, Mar. 2023, doi: 10.48550/ARXIV.2303.03017.
- [22] E. Riegler, G. E. Kerkelund, C. N. Manchon, M. A. Badiu, and B. H. Fleury, "Merging belief propagation and the mean field approximation: A free energy approach," *IEEE Trans. Inf. Theory*, Jan. 2013, doi: 10.1109/tit.2012.2218573.
- [23] O.-E. Barbu, C. Navarro Manchón, C. Rom, T. Balercia, and B. H. Fleury, "OFDM receiver for fast time-varying channels using block-sparse Bayesian learning," *IEEE Trans. Veh. Technol.*, Dec. 2016, doi: 10.1109/TVT.2016.2554611.
- [24] R. Prasad, C. R. Murthy, and B. D. Rao, "Joint channel estimation and data detection in MIMO-OFDM systems: A sparse Bayesian learning approach," *IEEE Trans. Signal Process.*, Oct. 2015, doi: 10.1109/TSP.2015.2451071.

- [25] G. E. Kirkelund, C. N. Manchon, L. P. B. Christensen, E. Riegler, and B. H. Fleury, "Variational message-passing for joint channel estimation and decoding in MIMO-OFDM," in *2010 IEEE Global Telecommun. Conf.*, Dec. 6–10, 2010, doi: 10.1109/GLOCOM.2010.5683839.
- [26] B. Jorgensen, *Statistical properties of the generalized inverse Gaussian distribution*, ser. Lecture notes in statistics. New York, NY, USA: Springer-Verlag, 1982, doi: <https://doi.org/10.1007/978-1-4612-5698-4>.
- [27] W. Rudin, *Principles of mathematical analysis*, 3rd ed. New York, NY, USA: McGraw-Hill, 1976.
- [28] A. Edelman and H. Murakami, "Polynomial roots from companion matrix eigenvalues," *Math. Comput.*, vol. 64, no. 210, pp. 763–776, 1995.
- [29] K. B. Petersen and M. S. Pedersen, *The Matrix Cookbook*. Technical University of Denmark, 2012, version Nov. 15, 2012, Accessed: Oct. 2023. [Online]. Available: <https://www.math.uwaterloo.ca/~hwolkowi/matrixcookbook.pdf>
- [30] M.-A. Badiu, T. L. Hansen, and B. H. Fleury, "Variational Bayesian inference of line spectra," *IEEE Trans. Signal Process.*, May 2017, doi: 10.1109/TSP.2017.2655489.
- [31] T. L. Hansen, B. H. Fleury, and B. D. Rao, "Superfast line spectral estimation," *IEEE Trans. Signal Process.*, Feb. 2018, doi: 10.1109/TSP.2018.2807417.
- [32] E. Leitinger, S. Grebien, B. Fleury, and K. Witrisal, "Detection and estimation of a spectral line in MIMO systems," in *2020 54th Asilomar Conf. Signals, Syst. and Computers*, Nov. 01–04, 2020, doi: 10.1109/IEEECONF51394.2020.9443381.
- [33] S. Grebien, E. Leitinger, K. Witrisal, and B. H. Fleury, "Super-resolution estimation of UWB channels including the diffuse component – an SBL-inspired approach," *ArXiv e-prints*, Aug. 2023, doi: 10.48550/arXiv.2308.01702.
- [34] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, "A consistent metric for performance evaluation of multi-object filters," *IEEE Trans. Signal Process.*, Aug. 2008, doi: 10.1109/TSP.2008.920469.
- [35] Y. Park, F. Meyer, and P. Gerstoft, "Graph-based sequential beamforming," *J. Acoustical Soc. Amer.*, Jan. 2023, doi: 10.1121/10.0016876.
- [36] X. Li, E. Leitinger, A. Venus, and F. Tufvesson, "Sequential detection and estimation of multipath channel parameters using belief propagation," *IEEE Trans. Wireless Commun.*, Oct. 2022, doi: 10.1109/TWC.2022.3165856.



Variational Inference of Structured Line Spectra Exploiting Group-Sparsity

Jakob Möderl, Franz Pernkopf, Klaus Witrisal, and Erik Leitinger

*Submitted to IEEE Transactions on Signal Processing, Feb. 2024,
doi: 10.48550/arXiv.2303.03017*

Abstract—This paper presents a method of decomposing a signal into several groups of related spectral lines. The frequencies of the spectral lines in each group are related to a parameter common to all spectral lines within the same group, e.g. the fundamental frequency of a harmonic series of spectral lines. The parameters of each group are estimated on a continuum by the proposed variational expectation-maximization (EM) algorithm. Furthermore, the number of groups and the number of spectral lines within each group are estimated via the group-sparse solutions obtained by estimating the latent variables in a hierarchical Bernoulli-Gamma-Gaussian prior model inspired by sparse Bayesian learning (SBL).

The performance of the proposed algorithm is demonstrated on three tasks: multi-pitch estimation, detection of extended objects using radar signals, and variational mode decomposition (VMD). On the Bach 10 dataset, which contains recordings of ten musical pieces, the proposed algorithm outperforms state-of-the-art model-based and machine learning-based multi-pitch estimation algorithms in terms of the fundamental frequency, i.e. pitch, detection accuracy. In addition, the extended object detection task shows how incorporating knowledge of the structural relationships between spectral lines into the estimation procedure can lead to performance gains compared to assuming independent spectral lines, especially under low signal-to-noise ratio (SNR) conditions. Finally, the VMD task is included to further demonstrate the versatility of the proposed algorithm.

This research was partly funded by the Austrian Research Promotion Agency (FFG) within the project SEAMAL Front (project number: 880598). Furthermore, the financial support by the Christian Doppler Research Association, the Austrian Federal Ministry for Digital and Economic Affairs and the National Foundation for Research, Technology and Development is gratefully acknowledged.

E.1 Introduction

Line spectral estimation (LSE) [1], i.e. estimating the frequencies and amplitudes of a superposition of complex exponential functions from a noisy measurement vector, is an ubiquitous task in signal processing. Solutions to this problem are applicable in many areas of physics and engineering, including range and direction estimation in radar and sonar, speech and music analysis, wireless channel estimation, molecular dynamics and geophysical exploration [1]. In many applications, the spectral lines can be organized into groups that have certain parameters in common. In speech or music analysis, the signal of each speaker during voiced speech or each tone of an instrument exhibits a harmonic structure with spectral lines at integer multiples of some base frequency [2–5]. In radar signal processing, extended objects give rise to multiple related target signals [6, 7]. These related target signals result in multiple correlated lines. Additionally, the signal models underlying many other signal processing tasks, e.g. variational mode decomposition (VMD) [8], can be approximated by a structured line spectrum. Typically, the number of groups as well as the group structure, i.e. which spectral line belongs to which group, is not known a priori and has to be estimated as well, significantly complicating the estimation process.

E.1.1 State of the Art

Common solutions to the LSE problem assume the number of spectral lines (i.e. the model order) to be known and no relation exists between the spectral lines. Examples thereof include subspace based methods such as MUSIC [9] or ESPRIT [10] as well as the maximum likelihood (ML) method [11, 12]. If the model order is not known, traditional methods combining detection and estimation select among multiple candidate models the one that optimizes a so-called information criterion, such as the Akaike or Bayesian information criterion, and the minimum description length. Each candidate model corresponds to a specific hypothesis on the number of spectral lines. See [1, 13] and references therein. However, using information-based approaches can be computationally expensive because a solution must be obtained for each considered model order before a specific solution is chosen.

Sparse signal reconstruction methods aim to reconstruct a signal based on a large dictionary matrix which is weighted with a sparse amplitude vector. Thereby, the model order is estimated as part of the process, alleviating the issue. A prominent instance of dictionary based sparse signal reconstruction method is the least absolute shrinkage and selection operator (LASSO) [14], which is also called basis pursuit denoising [15]. Further methods include matching pursuit [16], sparse Bayesian learning (SBL) [17–19], and SPICE [20]. Many of these algorithms have been extended to include a group structure, such as the group-LASSO [21–24], blockwise sparse regression [25], block matching pursuit [26], group-SBL [27–32], pattern-coupled SBL [33] and group-SPICE [34]. See [35–37] for detailed discussions about the similarities and differences of many of these methods. A disadvantage of using a fixed dictionary matrix is the spectral leakage induced by the model mismatch, which decreases the estimation performance [38, 39]. Thus, parametrized approaches have been developed such as the variational LSE (Valse) algorithm [40], the gridless-SPICE algorithm [41] and extensions of SBL to a continuous (i.e., infinite)

dictionary matrix with super-resolution capability¹ [42–45]. To the best of our knowledge, no SBL-based algorithm with super-resolution capabilities considers a group-structure in the dictionary matrix so far.

Methods to solve the LSE problem using a grouped approach can be found in the application of multi-pitch estimation, such as harmonic extensions for the Capon beamformer and the MUSIC principle, as well as an expectation-maximization (EM)-based estimator. See [4] for a collection of these methods. A more recent approach is based on block sparsity given a grid of fundamental frequencies [46]. However, since this approach is based on a fixed frequency grid, it suffers the same drawbacks as other sparse signal reconstruction methods with fixed dictionary matrices. To alleviate this issue, [47] proposes a block-sparse method for harmonic LSE based on a grouped continuous (infinite) dictionary matrix that is iteratively refined. Reference [47] is based on the superresolution iterative reweighted method [48], which belongs to the class of Type-I methods [36]. However, [36] showed that Type-I methods are a limiting case of Type-II methods. Furthermore, [36] showed that Type-II methods, such as SBL, possess several advantages compared to Type-I methods.

E.1.2 Contribution

We present an extension and generalization of SBL-based superresolution methods to infer structured line spectra. The developed algorithm outperforms related state-of-the art algorithms, e.g. [47], in the task of multi-pitch estimation on simulated and real data. Furthermore, our approach outperforms state-of-the art methods for (unstructured) LSE [44, 45, 49] when detecting and estimating extended objects in low-signal-to-noise ratio (SNR) conditions using radar signals and the VMD algorithm [8] when decomposing a signal consisting of two “intrinsic modes”. Specifically, the contributions of this work are as follows.

- We combine the Bernoulli-Gaussian model of VALSE [40] with the Gamma-Gaussian model used in SBL [18] to a layered hierarchical Bernoulli-Gamma-Gaussian model. The latent variables of this model are estimated using a variational-EM approach, resulting in a novel algorithm that estimates the continuous and discrete dictionary parameters while obtaining a solution that is sparse on two levels. (i) the number of groups and (ii) the number of spectral lines within each group.
- Different structural relations between the spectral lines in the model can be considered in our formulation. Thus, the developed algorithm can be applied to a variety of inference problems as illustrated in Section E.5.
- We derive the relation between the threshold governing the sparsity of groups and the threshold governing the sparsity of spectral lines within a group. This simplifies the process of tuning these thresholds to the application at hand.
- We demonstrate performance advantages on three different inference problems—multi-pitch estimation, detecting and estimating extended objects using radar signals and, VMD—using simulated data.

¹We define super-resolution as the ability of an algorithm to estimate the frequencies of the spectral lines on a continuum instead of restricting the estimates to a predefined grid.

- We evaluate the performance of the proposed algorithm on real multi-pitch data by applying it to the publicly available Bach 10 dataset containing recordings of a quartet playing 10 chorales by J. S. Bach [50]. Our algorithm achieves superior benchmark values of accuracy = 0.56, compared to accuracy = 0.47 and accuracy = 0.52, for the state-of-the-art model-based and pretrained methods of [47] and [51], respectively.

E.2 Signal Model and Bayesian Formulation

E.2.1 Signal Model

We consider an length- N signal vector $\mathbf{x} = [x(-\frac{N}{2}T_s) \ x((-\frac{N}{2} + 1)T_s) \cdots x((\frac{N}{2} - 1)T_s)]^T \in \mathbb{C}^N$, which contains the values of some continuous function $x(t)$ sampled at instances $\mathbf{t} = [-\frac{N}{2}T_s \ (-\frac{N}{2} + 1)T_s \cdots (\frac{N}{2} - 1)T_s]^T$ with regular sampling interval T_s . We assume that \mathbf{x} is a linear combination of spectral lines in noise $\boldsymbol{\epsilon}$, and the spectral lines can be structured into K groups as

$$\mathbf{x} = \sum_{k=1}^K \sum_{l \in \mathcal{S}_k} \alpha_{k,l} \boldsymbol{\psi}(\theta_k, l) + \boldsymbol{\epsilon}. \quad (\text{E.1})$$

Each group consists of one or multiple spectral lines $\boldsymbol{\psi}(\theta_k, l) = e^{j2\pi f(\theta_k, l)\mathbf{t}}$, also referred to as components, with frequencies $f(\theta_k, l)$.² The function $f(\theta_k, l)$ relates the frequency of the l th line in the k th group to the group parameter θ_k .³ Each spectral line $\boldsymbol{\psi}(\theta_k, l)$ is weighted with an amplitude $\alpha_{k,l} \in \mathbb{C}$ and the signal is corrupted by additive white Gaussian noise (AWGN) $\boldsymbol{\epsilon}$. We assume $\boldsymbol{\epsilon}$ follows a circular-symmetric complex Gaussian distribution, i.e., $p(\boldsymbol{\epsilon}) = \text{CN}(\boldsymbol{\epsilon}; \mathbf{0}, \lambda^{-1}\mathbf{I})$, with unknown precision λ .⁴

The main focus of the paper is the estimation of the group parameters θ_k together with the number of groups K and the sets of active components of each group \mathcal{S}_k . The amplitudes $\alpha_{k,l}$ and noise precision λ are considered nuisance parameters. Note that the signal model in (E.1) can be straightforwardly extended to multiple measurement vectors such as signals from a microphone or antenna array and to vector parameters as for example estimating multi-dimensional multipath parameters such as the angle-of-arrival, angle-of-departure, Doppler, and delay.

² $e^{\mathbf{a}}$ with $\mathbf{a} = [a_1 \ a_2 \ \cdots \ a_N]^T \in \mathbb{C}^N$ is defined to be a vector $e^{\mathbf{a}} \triangleq [e^{a_1} \ e^{a_2} \ \cdots \ e^{a_N}]^T$.

³Consider the audio spectrum resulting from multiple musical instruments playing several (constant) musical tones at the same time. Such a spectrum can be modeled by (E.1) with $f(\theta_k, l) = \theta_k l$ where θ_k is the fundamental frequency of the k th tone and \mathcal{S}_k the set of harmonics produced by the instrument playing this tone. Additional examples how to apply (E.1) to several practical problems are found in Section E.5.

⁴We denote the complex Gaussian PDF of the variable $\mathbf{x} \in \mathbb{C}^N$ with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$ as $\text{CN}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) = |\pi\boldsymbol{\Sigma}|^{-1} \exp\{-\frac{1}{\pi}(\mathbf{x} - \boldsymbol{\mu})^H \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})\}$, where $|\cdot|$ denotes the matrix determinant. Furthermore, we assume that $\mathbf{x} - \boldsymbol{\mu}$ is proper for all complex Gaussian random variables \mathbf{x} with mean $\boldsymbol{\mu}$.

E.2.2 Probabilistic Model

For inference, we define $\mathcal{S}_{\max} = \{l \mid L_{\min} \leq l \leq L_{\max}\}$ as the set of all potential components of a group and K_{\max} the maximum number of groups,⁵ and rewrite (E.1) as product of a large parametrized dictionary matrix $\Psi(\theta)$ whose columns contain all possible components of a large number of groups multiplied with a sparse amplitude vector α given by

$$\mathbf{x} = \Psi(\theta)\alpha + \epsilon \quad (\text{E.2})$$

where $\theta = [\theta_1 \ \theta_2 \ \cdots \ \theta_{K_{\max}}]^T$ is a vector of the corresponding group parameters θ_k and $\Psi(\theta_k) = [\psi(\theta_k, L_{\min}) \ \psi(\theta_k, L_{\min} + 1) \ \cdots \ \psi(\theta_k, L_{\max})]$ is a matrix whose columns contain all spectral lines parametrized by θ_k , and $\Psi(\theta) = [\Psi(\theta_1) \ \Psi(\theta_2) \ \cdots \ \Psi(\theta_{K_{\max}})]$. The sparse vector of all amplitudes is given by $\alpha = [\alpha_1^T \ \alpha_2^T \ \cdots \ \alpha_{K_{\max}}^T]^T$, where the amplitude vector for each possible group is given by $\alpha_k = [\alpha_{k, L_{\min}} \ \alpha_{k, L_{\min}+1} \ \cdots \ \alpha_{k, L_{\max}}]^T$.

To achieve sparsity on both levels, in the number of groups as well as in the number of components in each group, we propose to use a Bernoulli-Gamma-Gaussian prior model depicted in Figure E.1. We model the existence of each group with independent Bernoulli distributed random variables $z_k \in \{0, 1\}$, while simultaneously modeling the prior variance of each amplitude with a Gamma distributed random variable $\gamma_{k,l} \in \mathbb{R}_{\geq 0}$. Our model differs from [43, 45] in the addition of the Bernoulli-prior which is shown to increase resilience against the insertion of artificial components [40], while it differs from [40, 44] by using the Bernoulli-prior to model the existence of groups of several components instead of individual components. Note that we can constrain each group to contain at most a single spectral line with frequency θ_k by setting $f(\theta_k, l) = \theta_k$ and $\mathcal{S}_{\max} = \{0\}$. In this case, the hierarchical model is identical to [44]. Thus, the proposed method extends and generalizes the method in [44] to the estimation of structured line spectra. It differs from [44] by being based on the variational EM approach instead of maximizing a marginalized likelihood function.

For each potential group $k \in \{1, 2, \dots, K_{\max}\}$, we introduce latent binary random variables $\mathbf{z} = [z_1 \ z_2 \ \cdots \ z_{K_{\max}}]^T$, $z_k \in \{0, 1\}$ which select whether the k th group is active or not. If the k th group is not active, all amplitudes are zero, i.e. $\alpha_{k,l} = 0 \forall l$. The prior probability density function (PDF) of the amplitudes $\alpha_{k,l}$ of all active groups is further modeled by independent complex Gaussian PDFs with prior precision $\gamma_{k,l}$. Each precision $\gamma_{k,l}$ is again treated as latent random variable [17, 18, 43–45]. Thus, the prior PDF for an individual amplitude $\alpha_{k,l}$ conditioned on z_k and $\gamma_{k,l}$ is given by

$$p(\alpha_{k,l} | \gamma_{k,l}, z_k) = z_k \text{CN}(\alpha_{k,l}; 0, \gamma_{k,l}^{-1}) + (1 - z_k) \delta(\alpha_{k,l}) \quad (\text{E.3})$$

where $\delta(\cdot)$ is the Dirac delta distribution.⁶ The latent variables z_k represent sparsity-inducing priors on the group level and their priors are modeled as independent

⁵Since we can never expect to estimate more parameters than the number of observations, the actual values of L_{\min} , L_{\max} and K_{\max} do not impact the estimation performance of the proposed algorithm as long as they are chosen sufficiently large.

⁶This type of prior is also referred to as a “spike and slab” prior.

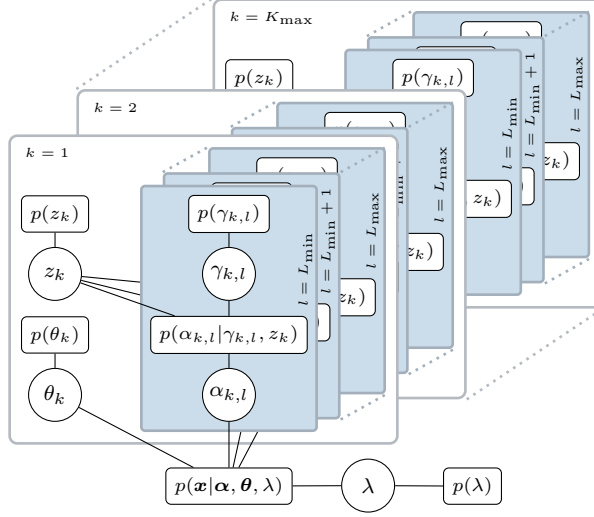


Figure E.1: Factor graph of the Bernoulli-Gamma-Gaussian model for sparse group estimates. Groups are indicated by a gray frame and components of groups in blue. Sparsity of components is promoted by the hyperparameters $\gamma_{k,l}$, while sparsity in groups is promoted by the hyperparameters z_k .

Bernoulli PDFs with \bar{z} as probability for $z_k = 1$, i.e.,

$$p(\mathbf{z}) = \prod_{k=1}^{K_{\max}} \bar{z}^{z_k} (1 - \bar{z})^{1-z_k}. \quad (\text{E.4})$$

Each precision $\gamma_{k,l}$ is modeled by independent Gamma PDFs $p(\gamma_{k,l}) = \text{Ga}(\gamma_{k,l}; \eta, \nu)$ with shape η and rate ν .⁷ Note that since the PDFs of the amplitude's precisions $p(\gamma_{k,l})$ are sparsity inducing, the according hierarchical model leads to many component amplitudes having a prior precisions of infinity, resulting in them being “removed from the model” as the corresponding amplitude is estimated to be zero as well [17–19]. The number of groups K and set of active components of each group \mathcal{S}_k are indirectly estimated by estimating z_k and $\gamma_{k,l}$, respectively.

Assuming independent priors $p(\alpha_{k,l}|\gamma_{k,l}, z_k)$, the joint prior PDF of the amplitudes is given by

$$p(\boldsymbol{\alpha}|\boldsymbol{\gamma}, \mathbf{z}) = \text{CN}(\boldsymbol{\alpha}_{\mathcal{S}}; \mathbf{0}, \boldsymbol{\Gamma}^{-1}) \prod_{\alpha_{k,l} \notin \boldsymbol{\alpha}_{\mathcal{S}}} \delta(\alpha_{k,l}). \quad (\text{E.5})$$

where $\boldsymbol{\gamma} = [\boldsymbol{\gamma}_1^T \boldsymbol{\gamma}_2^T \cdots \boldsymbol{\gamma}_{K_{\max}}^T]^T$ is obtained by stacking the vectors

$$\boldsymbol{\gamma}_k = [\gamma_{k,L_{\min}} \gamma_{k,L_{\min}+1} \cdots \gamma_{k,L_{\max}}]^T$$

of precisions $\gamma_{k,l}$ corresponding to the amplitudes α_k of the k th group, $\mathcal{S} = \mathcal{S}(\boldsymbol{\gamma}, \mathbf{z})$ is an index set such that $\boldsymbol{\alpha}_{\mathcal{S}}$ contains all nonzero elements of $\boldsymbol{\alpha}$ and $\boldsymbol{\gamma}_{\mathcal{S}}$ contains

⁷We denote the Gamma PDF with shape parameter a and rate parameter b as $\text{Ga}(x; a, b) = \frac{b^a}{\Gamma(a)} x^{a-1} e^{-bx}$, where $\Gamma(\cdot)$ is the gamma function.

the prior precisions corresponding the amplitudes α_S ,⁸ and $\mathbf{\Gamma} = \text{diag}(\gamma_S)$ is a diagonal matrix with the elements of the vector γ_S along the main diagonal. Furthermore, we assume the prior distribution of the noise variance to be a Gamma PDF $p(\lambda) = \text{Ga}(\lambda; \rho, \mu)$ with shape ρ and rate μ , since this is the conjugate prior for the precision of a Gaussian PDF.

Under the assumptions made in Section E.2, the likelihood $p(\mathbf{x}|\alpha, \theta, \lambda)$ is given by

$$p(\mathbf{x}|\alpha, \theta, \lambda) = \text{CN}(\mathbf{x}; \mathbf{\Psi}_S(\theta)\alpha_S, \lambda^{-1}\mathbf{I}). \quad (\text{E.6})$$

Introducing $p(\theta)$ as prior PDF of the parameters θ , the posterior PDF is proportional to

$$\begin{aligned} p(\alpha, \theta, \gamma, z, \lambda|\mathbf{x}) &\propto p(\mathbf{x}|\alpha, \theta, \lambda) \\ &\times p(\alpha|\gamma, z)p(\gamma)p(z)p(\theta)p(\lambda). \end{aligned} \quad (\text{E.7})$$

E.3 Variational Approximation

E.3.1 Mean-Field Factorization and Distribution Updates

To obtain a computationally tractable estimator, we apply a variational-EM approach [52] [53, Ch. 10]. We partition the set of variables $\Omega = \{\alpha, \theta, \gamma, z, \lambda\}$ into two subsets Ω_{VI} and Ω_{MAP} and aim to obtain (approximate) posterior PDFs for the variables $\Omega_{\text{VI}} \triangleq \{\alpha, \gamma, \lambda\}$, while maximum a-posteriori (MAP) point estimates $\hat{\theta}$ and \hat{z} are obtained for the remaining variables $\Omega_{\text{MAP}} \triangleq \{\theta, z\}$. Moreover, we apply a structured mean-field assumption to approximate the posterior PDF (E.7) at $\theta = \hat{\theta}$ and $z = \hat{z}$ with a proxy PDF $q(\alpha, \gamma, \lambda; \hat{\theta}, \hat{z})$ parameterized by $\hat{\theta}$ and \hat{z} . The proxy PDF q is postulated to factorize over the variables in Ω_{VI} as

$$q(\alpha, \gamma, \lambda; \hat{\theta}, \hat{z}) = q_\alpha(\alpha) q_\lambda(\lambda) \prod_{k=1}^{K_{\max}} \prod_{l=L_{\min}}^{L_{\max}} q_{\gamma,k,l}(\gamma_{k,l}) \quad (\text{E.8})$$

where the factors

$$q_j \in \mathcal{Q} \triangleq \{q_\alpha, q_\lambda, q_{\gamma,1,L_{\min}}, q_{\gamma,1,L_{\min}+1}, \dots, q_{\gamma,K_{\max},L_{\max}}\}$$

implicitly depend on $\hat{\theta}$ and \hat{z} . Besides the postulated factorization, we do not constrain the proxy PDF q . Thus, the shape of q_α , $q_{\gamma,k,l}$ and q_λ is determined by the variational optimization procedure. We minimize the Kullback-Leibler (KL) divergence from q to the posterior $p(\alpha, \theta, \gamma, z, \lambda|\mathbf{x})$ by maximizing the parametrized evidence lower bound (ELBO)

$$L(q; \hat{\theta}, \hat{z}) = \langle \ln p(\alpha, \hat{\theta}, \gamma, \hat{z}, \lambda|\mathbf{x}) - \ln q(\alpha, \gamma, \lambda; \hat{\theta}, \hat{z}) \rangle_q \quad (\text{E.9})$$

⁸We denote a vector subscripted by an index set α_S as the vector containing the elements of α whose indices are elements of S . Similarly, we denote for matrices $\mathbf{\Psi}_S(\theta)$ as the matrix formed by the columns of $\mathbf{\Psi}(\theta)$ whose indices are elements of S .

in an iterative fashion with respect to the factors of q and parameters $\hat{\boldsymbol{\theta}}$ and $\hat{\mathbf{z}}$ [52], [53, Ch. 10].⁹ We omit explicit iteration indices for brevity and refer to the last available estimates of the respective proxy PDFs and parameters.

Following the E-step of the variational-EM algorithm, the proxy PDF $q_j \in \mathcal{Q}$ that maximizes the ELBO, while keeping the remaining factors $q_i \in \mathcal{Q} \setminus \{q_j\}$ fixed, is found as

$$\ln q_j \stackrel{c}{\propto} \langle \ln p(\boldsymbol{\alpha}, \hat{\boldsymbol{\theta}}, \boldsymbol{\gamma}, \hat{\mathbf{z}}, \lambda | \mathbf{x}) \rangle_{\bar{q}_j} \quad (\text{E.10})$$

where $\bar{q}_j = \prod_{q_i \in \mathcal{Q} \setminus \{q_j\}} q_i$ denotes the product of all factors of the joint proxy PDF q except q_j and $\stackrel{c}{\propto}$ denotes that the right side is equal to the left side plus a constant, such that both sides are proportional to each other after applying the exponential function. However, only the factor q_λ relating to the noise precision is updated using this standard update as detailed in Appendix E.7.1. For the other factors of q and the parameters $\hat{\boldsymbol{\theta}}$ and $\hat{\mathbf{z}}$ we use either “marginalized” updates [40, 54] detailed in Section E.3.2 or “fast” updates [19, 32] detailed in Section E.3.3.

E.3.2 Joint Estimation of $\hat{\boldsymbol{\theta}}$, $\hat{\mathbf{z}}$ and q_α

There is a strong dependence between the proxy PDF q_α and the estimates $\hat{\mathbf{z}}$ and $\hat{\boldsymbol{\theta}}$, resulting in slow convergence if the usual M-step maximization [52], [53, Ch. 10] is used. To accelerate convergence, we use a similar approach as [40] and maximize the ELBO jointly with respect to q_α , $\hat{\mathbf{z}}$ and $\hat{\boldsymbol{\theta}}$. Let $q_\gamma = \prod_{k=1}^{K_{\max}} \prod_{l=L_{\min}}^{L_{\max}} q_{\gamma,k,l}$ be the product of all $q_{\gamma,k,l}$, we express the ELBO (E.9) as

$$\mathcal{L}(q_\alpha; \hat{\boldsymbol{\theta}}, \hat{\mathbf{z}}) = \text{const.} - \mathcal{D}_{\text{KL}}(q_\alpha \| t) + \ln Z(\hat{\boldsymbol{\theta}}, \hat{\mathbf{z}}) \quad (\text{E.11})$$

by ignoring all terms that do not depend on q_α , $\hat{\boldsymbol{\theta}}$ or $\hat{\mathbf{z}}$, and introducing the PDF

$$t(\boldsymbol{\alpha}; \hat{\boldsymbol{\theta}}, \hat{\mathbf{z}}) = \frac{1}{Z(\hat{\boldsymbol{\theta}}, \hat{\mathbf{z}})} \exp \langle \ln p(\boldsymbol{\alpha}, \hat{\boldsymbol{\theta}}, \boldsymbol{\gamma}, \hat{\mathbf{z}}, \lambda | \mathbf{x}) \rangle_{q_\lambda q_\gamma} \quad (\text{E.12})$$

where $\mathcal{D}_{\text{KL}}(q_\alpha \| t)$ denotes the KL divergence from q_α to t and $Z(\hat{\boldsymbol{\theta}}, \hat{\mathbf{z}})$ is a normalization constant. Since $\mathcal{D}_{\text{KL}} \geq 0$ with equality if, and only if, $t = q_\alpha$, (E.11) is maximized by

$$q_\alpha(\boldsymbol{\alpha}) = t(\boldsymbol{\alpha}; \hat{\boldsymbol{\theta}}, \hat{\mathbf{z}}) \quad (\text{E.13})$$

with estimates $\hat{\boldsymbol{\theta}}$ and $\hat{\mathbf{z}}$ determined by

$$\hat{\boldsymbol{\theta}}, \hat{\mathbf{z}} = \arg \max_{\boldsymbol{\theta}, \mathbf{z}} \ln Z(\boldsymbol{\theta}, \mathbf{z}). \quad (\text{E.14})$$

Let $\hat{\boldsymbol{\gamma}} = [\hat{\gamma}_{1,L_{\min}} \ \hat{\gamma}_{1,L_{\min}+1} \ \cdots \ \hat{\gamma}_{1,L_{\max}} \ \hat{\gamma}_{2,L_{\min}} \ \cdots \ \hat{\gamma}_{K,L_{\max}}]^T$ be the estimated mean value of $\boldsymbol{\gamma}$ based on the means $\hat{\gamma}_{k,l} = \langle \gamma_{k,l} \rangle_{q_{\gamma,k,l}}$, and $\hat{\mathcal{S}} = \mathcal{S}(\hat{\boldsymbol{\gamma}}, \hat{\mathbf{z}})$ be the current estimate of \mathcal{S} . As we derive in Appendix E.7.2, from (E.13) and (E.12) it follows that

$$q_\alpha(\boldsymbol{\alpha}) = \text{CN}(\boldsymbol{\alpha}_{\mathcal{S}}; \hat{\boldsymbol{\alpha}}, \hat{\mathbf{C}}) \prod_{\alpha_{k,l} \notin \boldsymbol{\alpha}_{\mathcal{S}}} \delta(\alpha_{k,l}) \quad (\text{E.15})$$

⁹We use $\langle f(a) \rangle_{q(a)}$ to denote the expectation of the function $f(a)$ with respect to some random variable a distributed according to the proxy PDF $q(a)$.

is a Gaussian PDF of all active amplitudes $\alpha_{\mathcal{S}}$ with the mean $\hat{\alpha}$ and covariance \hat{C} given by

$$\hat{\alpha} = \hat{\lambda} \hat{C} \hat{\Psi}^H \mathbf{x} \quad \hat{C} = (\hat{\lambda} \hat{\Psi}^H \hat{\Psi} + \hat{\Gamma})^{-1} \quad (\text{E.16})$$

where $\hat{\Psi} = \Psi_{\mathcal{S}}(\hat{\theta})$ is the matrix of all spectral lines parameterized by $\hat{\theta}$ with amplitudes estimated to be nonzero by $\hat{\gamma}$ and $\hat{\mathbf{z}}$, $\hat{\Gamma} = \text{diag}(\hat{\gamma}_{\mathcal{S}})$ is a diagonal matrix with the priors $\hat{\gamma}_{k,l}$ of the active components on its main diagonal and $\hat{\lambda} = \langle \lambda \rangle_{q_{\lambda}}$ is the estimated mean of λ .

To estimate $\hat{\theta}$ and $\hat{\mathbf{z}}$, we obtain $Z(\theta, \mathbf{z})$ from (E.12) as

$$Z(\theta, \mathbf{z}) = \int \exp \langle \ln p(\alpha, \theta, \gamma, \mathbf{z}, \lambda | \mathbf{x}) \rangle_{q_{\lambda} q_{\gamma}} d\alpha \quad (\text{E.17})$$

to express $\ln Z(\theta, \mathbf{z})$ in (E.14) as (see Appendix E.7.3)

$$\begin{aligned} \ln Z(\theta, \mathbf{z}) &\stackrel{e}{\propto} \hat{\lambda}^2 \mathbf{x}^H \Psi_{\mathcal{S}} C_{\mathcal{S}} \Psi_{\mathcal{S}}^H \mathbf{x} + \ln |C_{\mathcal{S}}| + \ln p(\theta) \\ &\quad + \ln p(\mathbf{z}) + \sum_{\hat{\gamma}_{k,l} \in \hat{\gamma}_{\mathcal{S}}} \langle \ln \gamma_{k,l} \rangle_{q_{\gamma,k,l}} \end{aligned} \quad (\text{E.18})$$

where $\tilde{\mathcal{S}} = \mathcal{S}(\hat{\gamma}, \mathbf{z})$, $\Psi_{\tilde{\mathcal{S}}} = \Psi_{\tilde{\mathcal{S}}}(\theta)$, $\hat{\Gamma}_{\tilde{\mathcal{S}}} = \text{diag}(\hat{\gamma}_{\tilde{\mathcal{S}}})$ and $C_{\tilde{\mathcal{S}}} = (\hat{\lambda} \Psi_{\tilde{\mathcal{S}}}^H \Psi_{\tilde{\mathcal{S}}} + \hat{\Gamma}_{\tilde{\mathcal{S}}})^{-1}$. Equation (E.18) is similar to the Type-II cost function that is obtained by maximizing the marginalized likelihood [43]. By plugging (E.18) into (E.14), $\hat{\theta}$ and $\hat{\mathbf{z}}$ are obtained by finding the global maximum. However, finding the global maximum jointly over θ and \mathbf{z} is computationally prohibitive. Thus, we explicitly express the dependence on a single pair of parameters $\{\theta_k, z_k\}$ and maximize (E.18) by coordinate ascent.

Let Ψ_k denote the columns of $\Psi_{\mathcal{S}}$ which correspond to the k th group and let the index \bar{k} refer to all the columns of matrices, or elements of vectors, which do not correspond to the k th group. Without loss of generality, we can reorder $\hat{\gamma}_{\tilde{\mathcal{S}}} = [\hat{\gamma}_{\bar{k}}^T \hat{\gamma}_k^T]^T$, $\theta = [\theta_{\bar{k}}^T \theta_k]^T$, $\mathbf{z} = [z_{\bar{k}}^T z_k]^T$, and $\Psi_{\tilde{\mathcal{S}}} = [\Psi_{\bar{k}} \Psi_k]$ such that the elements corresponding to the k th group are moved to the end. Assuming an uninformative prior for the parameters θ , i.e. $p(\theta) = \text{const.}$, we express the difference between $\ln Z$ with the k th group removed and $\ln Z$ including the k th group as

$$\begin{aligned} \Delta_k(\theta_k) &= \ln Z([\hat{\theta}_{\bar{k}}^T \theta_k]^T, [\hat{\mathbf{z}}_{\bar{k}}^T 1]^T) - \ln Z([\hat{\theta}_{\bar{k}}^T \theta_k]^T, [\hat{\mathbf{z}}_{\bar{k}}^T 0]^T) \\ &= \mathbf{u}^H C_k^{-1} \mathbf{u} + \ln |C_k| + \ln \frac{\bar{z}}{1 - \bar{z}} + \sum_{l \in \hat{\mathcal{S}}_k} (\chi_0 + \ln \hat{\gamma}_{k,l}) \end{aligned} \quad (\text{E.19})$$

where $\mathbf{u} = \hat{\lambda} C_k \Psi_k^H (\mathbf{I} - \hat{\lambda} \Psi_{\bar{k}} C_{\bar{k}} \Psi_{\bar{k}}^H) \mathbf{x}$, $C_k = (\hat{\lambda} \Psi_k^H \Psi_k + \hat{\Gamma}_k - \hat{\lambda}^2 \Psi_k^H \Psi_{\bar{k}} C_{\bar{k}} \Psi_{\bar{k}}^H \Psi_k)^{-1}$, $C_{\bar{k}} = (\hat{\lambda} \Psi_{\bar{k}}^H \Psi_{\bar{k}} + \hat{\Gamma}_{\bar{k}})^{-1}$, $\hat{\Gamma}_k = \text{diag}(\hat{\gamma}_k)$, $\hat{\Gamma}_{\bar{k}} = \text{diag}(\hat{\gamma}_{\bar{k}})$, $\hat{\mathcal{S}}_k = \{l | \hat{\gamma}_{k,l} < \infty\}$, and $\chi_0 = \text{digamma}(\eta + 1)$ is the digamma function evaluated at $\eta + 1$.

After $\hat{\theta}_k = \arg \max_{\theta_k} \Delta_k(\theta_k)$ is determined, we activate the k th group, i.e., $\hat{z}_k = 1$, by checking if $\Delta_k(\hat{\theta}_k) > 0$, which indicates an increase in $\ln Z$, i.e., an increase in the ELBO. In this case, the corresponding group is updated by $\hat{\theta}_k$. If $\Delta_k(\hat{\theta}_k) < 0$, the k th group is deactivated, i.e., $\hat{z}_k = 0$.

E.3.3 Fast Update of Prior Precision proxy PDFs $q_{\gamma,k,l}$

For the update of $q_{\gamma,k,l}$, given that the group is active, i.e., $\hat{z}_k = 1$ (due to the Bernoulli-prior model), we insert (E.7) and (E.6) into (E.10). After applying the logarithm on both sides, some simple manipulations, and ignoring all terms that do not depend on $\gamma_{k,l}$, we obtain

$$\begin{aligned} \ln q_{\gamma,k,l}(\gamma_{k,l}) &\stackrel{c}{\propto} \ln |\pi \mathbf{\Gamma}^{-1}|^{-1} - \langle \boldsymbol{\alpha}_{\hat{\mathcal{S}}}^H \mathbf{\Gamma} \boldsymbol{\alpha}_{\hat{\mathcal{S}}} \rangle_{\bar{q}_{k,l}} \\ &\quad + (\eta - 1) \ln \gamma_{k,l} - \gamma_{k,l} \nu \end{aligned} \quad (\text{E.20})$$

where $\bar{q}_{k,l}$ is the product of all the factors $q_j \in \mathcal{Q} \setminus \{q_{\gamma,k,l}\}$. Since $\mathbf{\Gamma} = \text{diag}(\boldsymbol{\gamma}_{\hat{\mathcal{S}}})$, we can simplify $\ln |\pi \mathbf{\Gamma}^{-1}|^{-1} = \ln \gamma_{k,l} + c_1$ and $\langle \boldsymbol{\alpha}_{\hat{\mathcal{S}}}^H \mathbf{\Gamma} \boldsymbol{\alpha}_{\hat{\mathcal{S}}} \rangle_{q_{\alpha}} = \gamma_{k,l} |\hat{\alpha}_{k,l}|^2 + |\hat{C}_{k,l}| + c_2$, where c_1 and c_2 are constants, $\hat{\alpha}_{k,l} = \langle \alpha_{k,l} \rangle_{q_{\alpha}}$, and $\hat{C}_{k,l}$ denotes the element on the main diagonal of $\hat{\mathbf{C}}$ that corresponds to the estimated variance of the amplitude $\alpha_{k,l}$. Thus, (E.20) is rewritten as a Gamma PDF given by

$$q_{\gamma,k,l}(\gamma_{k,l}) = \text{Ga}(\gamma_{k,l}; \eta + 1, \hat{M}_{\gamma,k,l}) \quad (\text{E.21})$$

with shape parameter $\eta + 1$ and rate parameter $\hat{M}_{\gamma,k,l} = |\hat{\alpha}_{k,l}|^2 + \hat{C}_{k,l} + \nu$. If the group is deactivated, i.e., $\hat{z}_k = 0$ (due to the Bernoulli-prior model), it follows that the proxy PDFs equal the prior PDFs, i.e., $q_{\gamma,k,l}(\gamma_{k,l}) = p(\gamma_{k,l})$. The mean $\hat{\gamma}_{k,l}$ of proxy PDF $q_{\gamma,k,l}$ is determined by

$$\hat{\gamma}_{k,l} = \begin{cases} \frac{\eta+1}{\hat{C}_{k,l} + \|\hat{\alpha}_{k,l}\|^2 + \nu} & \text{if } \hat{z}_k = 1 \\ \frac{\eta}{\nu} & \text{if } \hat{z}_k = 0. \end{cases} \quad (\text{E.22})$$

If the prior PDFs $p(\gamma_{k,l})$ are sparsity-inducing, many estimates $\hat{\gamma}_{k,l}$ will diverge, resulting in a sparse estimate for $\hat{\boldsymbol{\alpha}}$ [17, 19]. To obtain a fast convergence check, we consider Jeffery's prior, i.e., $p(\gamma_{k,l}) \propto \gamma_{k,l}^{-1}$ (if $\eta = \nu = 0$) and investigate the dependency of $\hat{\gamma}_{k,l}$ on $\hat{\boldsymbol{\alpha}}$ and $\hat{\mathbf{C}}$ [19]. Note that we can update q_{α} independently of $\hat{\boldsymbol{\theta}}$ and $\hat{\mathbf{z}}$ using (E.10), resulting in the same expression (E.15) for q_{α} . Thus, we express repeated cycles of updating q_{α} followed by updating $q_{\gamma,k,l}$ as a nonlinear map $G(\cdot)$. Inserting (E.16) into (E.22), each cycle i maps from the previous estimate of $\hat{\gamma}_{k,l}$ to the next as $\hat{\gamma}_{k,l}^{[i]} = G(\hat{\gamma}_{k,l}^{[i-1]})$. Hence, we can derive fast update rules for $\gamma_{k,l}$ by analyzing the stationary points of the map $G(\cdot)$. In line with [19], it can be shown that this map $G(\cdot)$ converges to

$$\hat{\gamma}_{k,l} \triangleq \hat{\gamma}_{k,l}^{[\infty]} = (|u_{k,l}|^2 - s_{k,l})^{-1} \quad \text{if } \frac{|u_{k,l}|^2}{s_{k,l}} > 1 \quad (\text{E.23})$$

and diverges otherwise. In (E.23), $s_{k,l}$ and $u_{k,l}$ are respectively given by

$$s_{k,l} = (\hat{\lambda} \boldsymbol{\psi}_{k,l}^H \boldsymbol{\psi}_{k,l} - \hat{\lambda}^2 \boldsymbol{\psi}_{k,l}^H \mathbf{\Psi}_{\bar{k},l} \mathbf{C}_{\bar{k},l} \mathbf{\Psi}_{\bar{k},l}^H \boldsymbol{\psi}_{k,l})^{-1}$$

and

$$u_{k,l} = s_{k,l} (\hat{\lambda} \boldsymbol{\psi}_{k,l}^H \mathbf{x} - \hat{\lambda}^2 \boldsymbol{\psi}_{k,l}^H \mathbf{\Psi}_{\bar{k},l} \mathbf{C}_{\bar{k},l} \mathbf{\Psi}_{\bar{k},l}^H \mathbf{x})$$

where $\boldsymbol{\psi}_{k,l} = \boldsymbol{\psi}(\hat{\theta}_k, l)$, $\mathbf{\Psi}_{\bar{k},l}$ is the dictionary matrix $\hat{\mathbf{\Psi}}$ with the column $\boldsymbol{\psi}_{k,l}$ removed, $\hat{\mathbf{\Gamma}}_{\bar{k},l}$ is a diagonal matrix containing the elements of $\hat{\boldsymbol{\gamma}}_{\hat{\mathcal{S}}}$ with $\hat{\gamma}_{k,l}$ removed, and

$\mathbf{C}_{k,l} = (\hat{\lambda} \Psi_{k,l}^H \Psi_{k,l} + \hat{\Gamma}_{k,l})^{-1}$. Thus, if (E.23) is fulfilled, we keep the l th component of the k th group in the model and discard it otherwise.

It can be shown that $\frac{|u_{k,l}|^2}{s_{k,l}}$ corresponds to the component SNR [19, 42, 55] and thus the condition $\frac{|u_{k,l}|^2}{s_{k,l}} > 1$ is equivalent to accepting any component that is only slightly above the noise level. However, this comes at the cost of accepting a significant number of false alarms. Let χ_1 be the minimum required component SNR, we can heuristically increase the threshold to

$$\frac{|u_{k,l}|^2}{s_{k,l}} > \chi_1 \geq 1 \quad (\text{E.24})$$

in order to reduce the false alarm rate at the cost of an increased missed detection rate. We refer the reader to [42, 55] for a theoretical derivation and more detailed analysis of the relationship between false alarm rate and threshold in the case of unstructured line spectra. Note that by increasing the threshold, we lose the guarantee that each update step will increase the ELBO and thus the convergence guarantee. However, in our simulations, we did not observe any noticeable negative impact on performance or convergence behavior by increasing the threshold.

E.3.4 Model Ambiguity and Constraints on Sparsity Parameters

The model (E.2) is ambiguous since many combinations of groups and active components within each group can lead to the same spectral lines. For example, if several lines are assigned to a group, one line can always be removed from the group and used to introduce a new group parameterized to produce the same line spectrum as the single group. However, this increases the degrees of freedom in the model. To avoid such overparameterization, the introduction of a new group into the model must be more costly than adding a new component within an existing group. To ensure this, we choose the group existence prior \bar{z} such that the activation of a group k based on $\Delta_k(\hat{\theta}_k)$ is stricter than the selection of a single spectral line based on (E.24) by using χ_1 . We start by using (E.19) to determine the activation rule of a group consisting only of a single component, which is given by

$$\frac{|u_{k,l}|^2}{s_{k,l}} \frac{1}{1 + \hat{\gamma}_{k,l} s_{k,l}} + \ln \frac{\hat{\gamma}_{k,l} s_{k,l}}{1 + \hat{\gamma}_{k,l} s_{k,l}} + \ln \frac{\bar{z}}{1 - \bar{z}} + \chi_0 > 0 \quad (\text{E.25})$$

using $\mathbf{C}_k = [(s_{k,l}^{-1} + \hat{\gamma}_{k,l})^{-1}] = [\frac{s_{k,l}}{1 + \hat{\gamma}_{k,l} s_{k,l}}]$ and $\mathbf{u} = [\frac{u_{k,l}}{1 + \hat{\gamma}_{k,l} s_{k,l}}]$. Note that (E.25) depends on the component SNR $\frac{|u_{k,l}|^2}{s_{k,l}}$, on the prior precision $\hat{\gamma}_{k,l}$, and the Bernoulli parameter \bar{z} . Using (E.25) and (E.24), we ensure that the inclusion of new groups with only a single component is penalized more than the inclusion of new components within a group by choosing the group existence prior \bar{z} such that

$$(1 + \hat{\gamma}_{k,l} s_{k,l}) \left(\ln \frac{1 + \hat{\gamma}_{k,l} s_{k,l}}{\hat{\gamma}_{k,l} s_{k,l}} + \ln \frac{1 - \bar{z}}{\bar{z}} - \chi_0 \right) > \chi_1 \quad (\text{E.26})$$

holds for any value of $\hat{\gamma}_{k,l} s_{k,l}$. Since $\hat{\gamma}_{k,l}$ and $s_{k,l}$ are both strictly positive quantities, we have $(1 + \hat{\gamma}_{k,l} s_{k,l}) > 1$ and $\ln \frac{1 + \hat{\gamma}_{k,l} s_{k,l}}{\hat{\gamma}_{k,l} s_{k,l}} > 0$. Thus, choosing

$$\bar{z} < \frac{1}{1 + \exp(\chi_0 + \chi_1)}. \quad (\text{E.27})$$

is sufficient for (E.26) to hold. Using (E.27), we express the cluster existence prior in terms of a threshold χ_2 , i.e. $\bar{z} = \frac{1}{1+\exp(\chi_0+\chi_2)}$ with $\chi_2 > \chi_1$.

E.4 Algorithm Implementation

Updating the parameters $\hat{\theta}$ as well as the proxy distributions $q_j \in \mathcal{Q}$ in the way described in the previous section will converge towards a local optimum of the ELBO. However, there might exist several local optima and the obtained solution depends on the initialization as well as the order in which the updates are performed. In this section, we define an iterative schedule for updating the factors q_α , q_λ and $q_{\gamma,k,l}$ and to estimate $\hat{\theta}$ and \hat{z} as well as an initialization. The resulting algorithm is outlined in Algorithm 2. We choose Jeffrey's priors ($\rho = \mu = \eta = \nu = 0$) for $p(\gamma_{k,l})$ and $p(\lambda)$, since these priors are non-informative for the noise precision and it allows us to use the fast convergence check developed in Section E.3.3 for the variances $\hat{\gamma}_{k,l}$.

Without loss of generality, we can reorder the groups such that $\hat{z} = [1 \cdots 1 \ 0 \cdots 0]$ is a vector of \hat{K} leading ones followed by $K_{\max} - \hat{K}$ zeros. Therefore, we only need to keep track of the estimated number of active groups \hat{K} and their parameters $\hat{\theta}_1$ through $\hat{\theta}_{\hat{K}}$ instead of the full vectors \hat{z} and $\hat{\theta}$. Similarly, instead of keeping the full vector $\hat{\gamma}$, we keep track only of the priors $\hat{\gamma}_{k,l} < \infty$ and denote their respective indices with index sets \hat{S}_k for all $k \in \{1, 2, \dots, \hat{K}\}$. We start with an empty model (bottom-up initialization), where $\hat{K} = 0$, $\hat{\theta}$, $\hat{\gamma}$, and $\hat{\alpha}$ are empty vectors, and \hat{C} is an empty matrix. The noise precision is initialized using the signal energy as $\hat{\lambda} = \frac{N}{\|\mathbf{x}\|^2}$. We then repeatedly alternate between searching for a new set of components to add to the model and updating the existing sets. We stop the procedure when the change in $\hat{\mathbf{x}} = \hat{\Psi}\hat{\alpha}$ from one iteration to the next is below a threshold and the search does not find a new group to add.

To search for a new group, we would ideally find a combination of group parameter $\hat{\theta}_{\hat{K}+1}$ and priors $\hat{\gamma}_{\hat{K}+1,l}$ (and thereby also of the group structure $\hat{S}_{\hat{K}+1}$) which maximizes $\Delta_{\hat{K}+1}$ in (E.19). However, since this requires an optimization over all possible subsets $\hat{S}_{\hat{K}+1} \subseteq \mathcal{S}_{\max}$ this is intractable. Thus, we choose a single component $l \in \mathcal{S}_{\max}$, e.g. $l = 1$, and consider a new group parametrized by $\hat{\theta}_{\hat{K}+1} = \arg \max_{\theta} |\psi^H(\theta, l)\hat{\mathbf{x}}_{\text{res}}|$, where $\hat{\mathbf{x}}_{\text{res}} = \mathbf{x} - \hat{\Psi}\hat{\alpha}$ is the residual signal. Next, we perform Algorithm 3 to find other related components in the proposed group and calculate the priors $\hat{\gamma}_{\hat{K}+1,l}$ for all l within a search region $\mathcal{S}_{\text{search}}$ discussed in detail in the next paragraph, i.e. for all $l \in \mathcal{S}_{\text{search}}$. Thereby we obtain an estimate of the group structure $\hat{S}_{\hat{K}+1} = \{l : \hat{\gamma}_{\hat{K}+1,l} < \infty\}$. Only after we have an estimate of the group parameter $\hat{\theta}_{\hat{K}+1}$ as well as the group structure encoded in $\hat{S}_{\hat{K}+1}$ of the new potential group, we check if adding the whole group to the model increases the ELBO by checking whether $\Delta_{\hat{K}+1}(\hat{\theta}_{\hat{K}+1}) > 0$. If this is the case, we add the group to the model. Otherwise, no new group is added.

After searching for a new group, we iterate over all groups $k \in \{1, 2, \dots, \hat{K}\}$ and for each one we first perform an E-step to update the distributions $q_{\gamma,k,l}$, followed by an M-step to update $\hat{\theta}_k$ and \hat{z}_k . Lastly, we update the amplitude and noise distributions q_α and q_λ . The update of the distributions $q_{\gamma,k,l}$ is outlined in Algorithm 3 and entails both updating the prior of existing components $\hat{\gamma}_{k,l}$ for all

Algorithm 2 Group-Sparse Super-Resolution LSE

Input: Signal vector \mathbf{x} , χ_1 , $\bar{z} = 1/(1 + \exp(\chi_0 + \chi_2))$.

Output: Model order \hat{K} , parameters $\hat{\boldsymbol{\theta}}$, and amplitudes $\hat{\boldsymbol{\alpha}}$.

Initialize $\hat{K} = 0$, $\hat{\lambda} = \frac{N}{\|\mathbf{x}\|^2}$, and $\hat{\boldsymbol{\alpha}}$, $\hat{\boldsymbol{\theta}}$, $\hat{\boldsymbol{\gamma}}$ as empty vectors.

repeat

$\hat{\mathbf{x}}_{\text{res}} \leftarrow \mathbf{x} - \hat{\Psi}\hat{\boldsymbol{\alpha}}$.

$\hat{\theta}_{\hat{K}+1} \leftarrow \arg \max_{\theta} |\boldsymbol{\psi}^H(\theta, 1)\hat{\mathbf{x}}_{\text{res}}|$.

 Execute Alg. 3 to estimate priors $\hat{\boldsymbol{\gamma}}_{\hat{K}+1}$ and $\hat{\mathcal{S}}_{\hat{K}+1}$.

 Compute $\Delta_{\hat{K}+1}(\hat{\theta}_{\hat{K}+1})$ from (E.19)

if $\Delta_{\hat{K}+1}(\hat{\theta}_{\hat{K}+1}) > 0$ **then**

$\hat{K} \leftarrow \hat{K} + 1$.

$\hat{\boldsymbol{\theta}} \leftarrow [\hat{\boldsymbol{\theta}}^T, \hat{\theta}_{\hat{K}+1}]^T$, $\hat{\boldsymbol{\gamma}} \leftarrow [\hat{\boldsymbol{\gamma}}^T, \hat{\boldsymbol{\gamma}}_{\hat{K}+1}^T]^T$.

end if

for all groups $k \in \{1, 2, \dots, \hat{K}\}$ **do**

 Execute Alg. 3 to update priors $\hat{\boldsymbol{\gamma}}_k$ and $\hat{\mathcal{S}}_k$.

 Find $\hat{\theta}_k = \arg \max_{\theta_k} \Delta_k(\theta_k)$ from (E.19).

 Compute $\Delta_k(\hat{\theta}_k)$ from (E.19).

if $\Delta_k(\hat{\theta}_k) \leq 0$ **then**

$\hat{K} \leftarrow \hat{K} - 1$, $\hat{\boldsymbol{\theta}} \leftarrow \hat{\boldsymbol{\theta}}_{\hat{K}}$, $\hat{\boldsymbol{\gamma}} \leftarrow \hat{\boldsymbol{\gamma}}_{\hat{K}}$.

end if

end for

 Compute $\hat{\boldsymbol{\alpha}}$ and $\hat{\mathbf{C}}$ from (E.16).

 Compute $\hat{\lambda}$ from (E.42).

until stopping criterion.

Algorithm 3 Fast update of priors $\hat{\gamma}_k$

Input: Signal vector \mathbf{x} , parameters $\hat{\boldsymbol{\theta}}, \hat{\gamma}, \hat{\mathcal{S}}_k, \hat{\lambda}$, index k , χ_1 .

Output: Prior precisions $\hat{\gamma}_k$ and $\hat{\mathcal{S}}_k$ of the k th group.

```

for all Components  $l \in \hat{\mathcal{S}}_k$  do
     $s_{k,l} \leftarrow (\hat{\lambda} \psi_{k,l}^H \psi_{k,l} - \hat{\lambda}^2 \psi_{k,l}^H \Psi_{k,l} C_{k,l} \Psi_{k,l}^H \psi_{k,l})^{-1}$ .
     $u_{k,l} \leftarrow s_{k,l} (\hat{\lambda} \psi_{k,l}^H \mathbf{x} - \hat{\lambda}^2 \psi_{k,l}^H \Psi_{k,l} C_{k,l} \Psi_{k,l}^H \mathbf{x})$ .
    if  $\frac{|u_{k,l}|^2}{s_{k,l}} > \chi_1$  then
         $\hat{\gamma}_{k,l} \leftarrow (|u_{k,l}|^2 - s_{k,l})^{-1}$ .
    else
        Remove  $\hat{\gamma}_{k,l}$  from  $\hat{\gamma}_k$ .
         $\hat{\mathcal{S}}_k \leftarrow \hat{\mathcal{S}}_k \setminus \{l\}$ .
    end if
end for
for all Components  $l \in \mathcal{S}_{\text{search}} \setminus \hat{\mathcal{S}}_k$  do
     $s_{k,l} \leftarrow (\hat{\lambda} \psi_{k,l}^H \psi_{k,l} - \hat{\lambda}^2 \psi_{k,l}^H \hat{\Psi} \hat{C} \hat{\Psi}^H \psi_{k,l})^{-1}$ .
     $u_{k,l} \leftarrow s_{k,l} (\hat{\lambda} \psi_{k,l}^H \mathbf{x} - \hat{\lambda}^2 \psi_{k,l}^H \hat{\Psi} \hat{C} \hat{\Psi}^H \mathbf{x})$ .
    if  $\frac{|u_{k,l}|^2}{s_{k,l}} > \chi_1$  then
         $\hat{\gamma}_k \leftarrow [\hat{\gamma}_k^T, (|u_{k,l}|^2 - s_{k,l})^{-1}]^T$ .
         $\hat{\mathcal{S}}_k \leftarrow \hat{\mathcal{S}}_k \cup \{l\}$ .
    end if
end for

```

$l \in \hat{\mathcal{S}}_k$ as well as looking for new components to add to the group. Intuitively, we would calculate $u_{k,l}$ and $s_{k,l}$ for all $l \in \mathcal{S}_{\text{max}}$ to check whether the component should be added or kept in the group or if it should be discarded. However, depending on the application and our choice of \mathcal{S}_{max} this could be suboptimal. Consider the case of extended object detection. Since we do not want to constrain the size of each object, we are encouraged to use a large range for \mathcal{S}_{max} . However, if two small objects are close to each other, this would potentially result in the estimation of only a single group covering both objects with a few spectral lines deactivated in the middle. To prevent this, we can constrain the search space to $\mathcal{S}_{\text{search}} \subseteq \mathcal{S}_{\text{max}}$ depending on the application. E.g. a reasonable choice for the example of extended object detection is to look for new components only in the neighbourhood of the currently existing ones by setting $\mathcal{S}_{\text{search}} = \{\min(\hat{\mathcal{S}}_k) - 1, \min(\hat{\mathcal{S}}_k), \dots, \max(\hat{\mathcal{S}}_k) + 1\} \cap \mathcal{S}_{\text{max}}$. If such a constrained search space is used, it can be beneficial to run a few updates of each group before adding a new group in order to explore the search space quicker and avoid introducing new groups for components which would be covered by another existing group anyway.

E.5 Applications and Results

E.5.1 Multi-Pitch Estimation

Multi-pitch estimation is a fundamental problem in audio signal processing [2–5, 46, 47]. The goal of multi-pitch estimation is to decompose the signal into several sources, each of which is modeled as a sum of harmonics, giving rise to the harmonically structured model

$$\mathbf{x} = \sum_{k=1}^K \sum_{l \in \mathcal{S}_k} \alpha_{k,l} e^{j2\pi l f_{0,k} t} + \boldsymbol{\epsilon}. \quad (\text{E.28})$$

Note that (E.28) is an instance of (E.1) since it represents a weighted mixture of multiple fundamental frequencies, i.e., pitches, $f_{0,k}$ and their harmonics given by $lf_{0,k}$ which can be expressed as $f(\theta_k, l) = \theta_k l$ with $\theta_k = f_{0,k}$. We aim to estimate the number of sources K along with the fundamental frequency $f_{0,k}$ of each source, the amplitudes $\alpha_{k,l}$, and the index set \mathcal{S}_k . Even though audio signals are typically real-valued, we can apply the complex-valued signal model by calculating the (down-sampled) discrete-time analytical signal [56].

To adapt the proposed algorithm to multi-pitch estimation, we refine the search strategy to fit the task at hand. When looking for new components, we consider all harmonics up to a relative frequency of $\theta_k l = 1$. Thus, we use $\mathcal{S}_{\text{search}} = \{1, 2, \dots, \text{floor}(\frac{1}{\hat{\theta}_k})\}$. To find the true fundamental frequency, we also perform a fractional search during which we search for components at fractions $l' \in \{\frac{1}{2}, \frac{1}{3}, \dots, \frac{1}{\text{floor}(N\hat{\theta}_k)}\}$ of the current estimate. If we find one such component we stop the fractional search and add l' to $\hat{\mathcal{S}}_k$. In order to obtain integer relations between all components we then re-parametrize $\hat{\theta}'_k = \hat{\theta}_k l'$ and $\hat{\mathcal{S}}'_k = \frac{\hat{\mathcal{S}}_k}{l'}$.

Numerical Analysis

To highlight the robustness of our algorithm against AWGN, we generate a signal of length $N = 100$ samples. The fundamental frequencies of $K = 2$ sources with 6 harmonics each are drawn uniformly from the interval $[0.025, 0.1]$. A minimum separation of the fundamental frequencies of 0.02 was ensured. All components have unit amplitude with a uniformly drawn phase, i.e. $\alpha_{k,l} = e^{j\varphi_{k,l}}$ with $\varphi_{k,l}$ drawn independently from the uniform PDF $p(\varphi_{k,l}) = \text{U}(\varphi_{k,l}, 0, 2\pi)$.¹⁰ In this experiment, we performed 1000 simulation runs. To evaluate the estimation accuracy and cardinality error of the estimated fundamental frequencies in a single metric, we use the optimal subpattern assignment (OSPA) metric [57], averaged over all simulation runs. We calculate the OSPA errors based on the Euclidean metric with cutoff parameter $\frac{2}{N} = 0.02$ and order $p = 1$.

For the proposed algorithm, we used thresholds for component and group sparsity of $\chi_1 = 7$ dB and $\chi_2 = 10$ dB, respectively. We compare our algorithm to the block super-resolution iteratively reweighted (BSURE-IR) algorithm¹¹ [47]

¹⁰We denote the uniform PDF of the variable x over the interval $[a, b]$ as $\text{U}(x; a, b) = \frac{1}{b-a}$ for $a \leq x \leq b$ and 0 otherwise.

¹¹https://www.maths.lu.se/fileadmin/maths/personal_staff/Andreas_Jakobsson/BSURE.zip




Figure E.2: OSPA of the estimated fundamental frequencies of a signal containing (a) 2 harmonic sources with fundamental frequencies drawn randomly from the interval $[0.025, 0.1]$ and (b) a major triad with fundamental frequencies f_0 , $\frac{5}{4}f_0$ and $\frac{3}{2}f_0$.

and the approximate nonlinear-least-squares (ANLS) EM algorithm¹² of [5]. The BSURE-IR algorithm was initialized with a grid of 15 frequency points in the interval $[0.025, 0.1]$ and the maximum allowed harmonic order was set to 6. The search interval for the ANLS-EM algorithm was set to $[0.025, 0.1]$ and an FFT size of 2^{12} was used.

Figure E.2a shows the mean OSPA for all three algorithms versus the SNR defined as $\text{SNR} = \frac{\|\mathbf{x} - \boldsymbol{\epsilon}\|^2}{\|\boldsymbol{\epsilon}\|^2}$. The BSURE-IR algorithm was not able to find any fundamental frequencies for SNRs < 10 dB, as indicated by the OSPA being equal to the cutoff-distance c . Even for high SNRs of 20 dB and more, the performance of BSURE-IR was worse than that of the proposed algorithm. The performance of the ANLS-EM algorithm is similar to the proposed algorithm. For low SNRs of 5 dB and less the ANLS-EM performs slightly better than the proposed algorithm in this example. However, both comparison algorithms have larger average runtime (averaged over simulation runs) than the proposed algorithm. The average runtime of the BSURE-IR algorithm strongly depends on the SNR. It is varying from 3.7 seconds for $\text{SNR} = 25$ dB to 20.0 seconds for $\text{SNR} = 5$ dB. In contrast, the average runtimes of the proposed algorithm and the ANLS-EM algorithm remain constant across SNR. The proposed method takes 0.1 seconds, while the ANLS-EM algorithm takes 5.2 seconds.

¹²<https://www.morganclaypool.com/page/multi-pitch>

In a second experiment, we generate a signal consisting of three notes with fundamental frequency f_0 , $\frac{5}{4}f_0$ and $\frac{3}{2}f_0$, i.e. a major triad, and $N = 100$ samples. Each note consists of six harmonics drawn randomly from $\{1, 2, \dots, 7\}$. The fundamental frequency of the base note f_0 was drawn uniformly from the interval $[0.07, 0.08]$ and all harmonics are generated with unit amplitude and uniform phase. We performed 1000 simulations runs. The BSURE-IR algorithm was initialized with 15 grid points in the interval $[0.06, 0.13]$ while the ANLS-EM algorithm was given the same range as frequency prior and the FFT-size was again set to 2^{12} . For the proposed algorithm, the same thresholds as in the previous experiment are used. Figure E.2b shows again the mean OSPA versus the SNR. In this scenario the proposed algorithm significantly outperforms the ANLS-EM algorithm for almost all SNR values. This is mainly due to the fact that the ANLS-EM algorithm underestimates the model order in most cases, even for high SNR values of 20 dB and more. This indicates a performance degradation of the ANLS-EM algorithm in more complex scenarios containing more fundamental frequencies and overlapping harmonics. Again, The performance of BSURE-IR is the worst of the three, even in the high SNR regime.

Real Music Signals

To evaluate the performance of the proposed algorithm on real data, we apply it to estimate the fundamental frequencies in the Bach 10 dataset¹³ [58]. The dataset contains 10 chorales of J.S. Bach played by a quartet consisting of a violin, a clarinet, a saxophone and a bassoon. Each piece lasts between 25-40 seconds with all instruments playing nearly all the time. The audio of each instrument was recorded individually while the musician listened to the others via headphones. The fundamental frequency of each instrument was extracted from the individual recordings using the YIN single-pitch estimator [59] as ground truth. Obvious errors in the ground truth are corrected manually. The audio signal was segmented into frames of 45 ms with a 10 ms stride between frames and the proposed algorithm was applied to each frame individually. Since the SNR and sampling rate of the audio signal is large, we select rather large thresholds of $\chi_1 = 15$ dB and $\chi_2 = 21$ dB, respectively. Furthermore, we constrain the fundamental frequencies to be within the interval (75 Hz, 10 kHz) to reduce halving errors. Each fundamental frequency in the ground truth was considered matched if an estimated fundamental frequency deviated from it no more than a half of a semitone, i.e. if it deviates no more than 3 % from the true fundamental frequency.

As comparison methods we used two model-based methods, BSURE-IR [47] and PEARLS [60], as well as a state-of-the-art machine learning method [51], abbreviated as BW15, that is pre-trained on the instruments in the dataset. The results for BW15 are take from [60] since they reported an overall better accuracy for the method compared to the original paper. The ANLS-EM was not included in the table, since preliminary investigations showed a significantly worse performance on the dataset than the other algorithms. To compare the algorithms, we use the common performance metrics accuracy, precision, recall, and the F_1 score. The

¹³<https://labsites.rochester.edu/air/resource.html>

Table E.1: Performance measures for the proposed algorithm evaluated on the Bach-10 dataset

Method	F ₁	Accuracy	Precision	Recall	Pre-Trained
Prop. Alg.	0.72	0.56	0.73	0.70	No
BW15	0.67	0.52	0.68	0.68	Yes
BSURE-IR	0.64	0.47	0.68	0.54	No
PEARLS	0.60	0.44	0.56	0.51	No

metrics are calculated by [61]

$$\text{Accuracy} = \frac{\text{TP}}{\text{TP} + \text{FP} + \text{FN}} \quad (\text{E.29})$$

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad (\text{E.30})$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (\text{E.31})$$

$$F_1 = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (\text{E.32})$$

where TP is the number of fundamental frequencies matched between the ground truth and the estimate in all frames, FP is the number of false positives, i.e., the number of estimated fundamental frequencies that did not match to any ground truth fundamental frequency, and FN is the number of ground truth fundamental frequencies that did not match any fundamental frequency estimate. Table E.1 lists the accuracy, precision, recall and the F₁ score of the proposed algorithm along with several comparison methods.

The algorithm proposed is capable of capturing most of the fundamental frequencies with very few false positives. It outperforms all three state-of-the-art comparison methods, including the pre-trained method BW15, which is based on the instruments in the dataset (see Table E.1). Figure E.3 shows an estimation-example of the proposed algorithm for several frames (15 seconds) of the choral “*Ach Gott und Herr*” from the dataset. In particular, it shows the estimated fundamental frequencies of the instruments along with the ground truth frequencies.

Note that in this work we considered only methods which work on the basis of individual frames to keep the comparison fair. There exist several methods for multi-pitch estimation which not only estimate the fundamental frequencies on a frame-by-frame basis but track notes over multiple frames. Thereby, these methods are able to achieve an F₁ score of up to 0.85 on the Bach 10 dataset. See e.g. [62] and references therein. However, these methods are usually data-driven whereas our algorithm is the best purely model-based algorithm. Since it does not rely on training data, we expect the proposed algorithm to generalize better to new datasets, out-of-domain data or noisy input data than data-driven approaches. Additionally, we expect that the performance of the proposed algorithm can be increased significantly by fusing information between frames in line with [63].

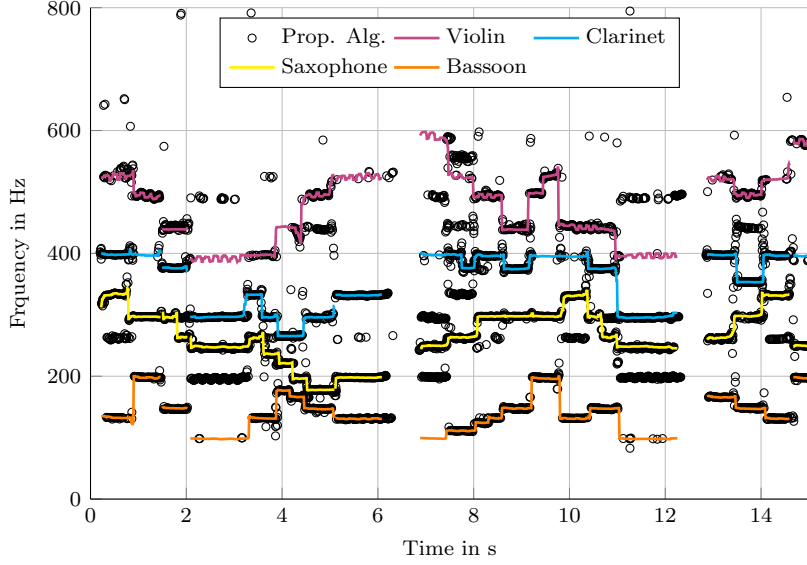


Figure E.3: Fundamental frequency estimates for 15 seconds of the choral *Ach Gott und Herr* compared to the ground truth (best viewed in color). The performance of the proposed algorithm for the depicted timespan is accuracy = 0.55, precision = 0.75 and recall = 0.68.

E.5.2 Extended Object Detection

A well-studied problem in radar signal processing is the detection of extended objects [6, 7]. An extended object is defined as a volume over which scatter points are distributed that correspond to a single physical target such as a car. We study a setup with a transmitter and a receiver, each equipped with a single antenna. A complex baseband signal $s(t)$ is transmitted and reflected by each of L_k scattering points of the k th target and reaches the receiving antenna with a delay $\tau_{k,l} > 0$ and complex amplitude $\alpha_{k,l} \in \mathbb{C}$. The corresponding received baseband signal is given by [7]

$$r(t) = \sum_{k=1}^K \sum_{l=1}^{L_k} \alpha_{k,l} s(t - \tau_{k,l}) + \epsilon(t). \quad (\text{E.33})$$

Let $\Delta f = \frac{1}{NT_s}$, $R(f)$ be the Fourier transform of $r(t)$ and $\mathbf{r} = [R(-\frac{N}{2}\Delta f) R((-\frac{N}{2} + 1)\Delta f) \cdots R((\frac{N}{2} - 1)\Delta f)]^T$. Similarly, let $S(f)$ be the Fourier transform of $s(t)$, $\mathbf{s} = [S(-\frac{N}{2}\Delta f) S((-\frac{N}{2} + 1)\Delta f) \cdots S((\frac{N}{2} - 1)\Delta f)]^T$, $\boldsymbol{\epsilon}$ be a sampled vector of the noise in frequency domain, and $\mathbf{f} = [-\frac{N}{2}\Delta f \ (-\frac{N}{2} + 1)\Delta f \cdots (\frac{N}{2} - 1)\Delta f]^T$. We can then express (E.33) in the frequency domain as

$$\mathbf{r} = \sum_{k=1}^K \sum_{l=1}^{L_k} \alpha_{k,l} \mathbf{s} \odot e^{-j2\pi\tau_{k,l}\mathbf{f}} + \boldsymbol{\epsilon} \quad (\text{E.34})$$

where \odot denotes the elementwise multiplication of two vectors. If the pulse is sufficiently short in time domain, we can apply the sampling theorem to approximate

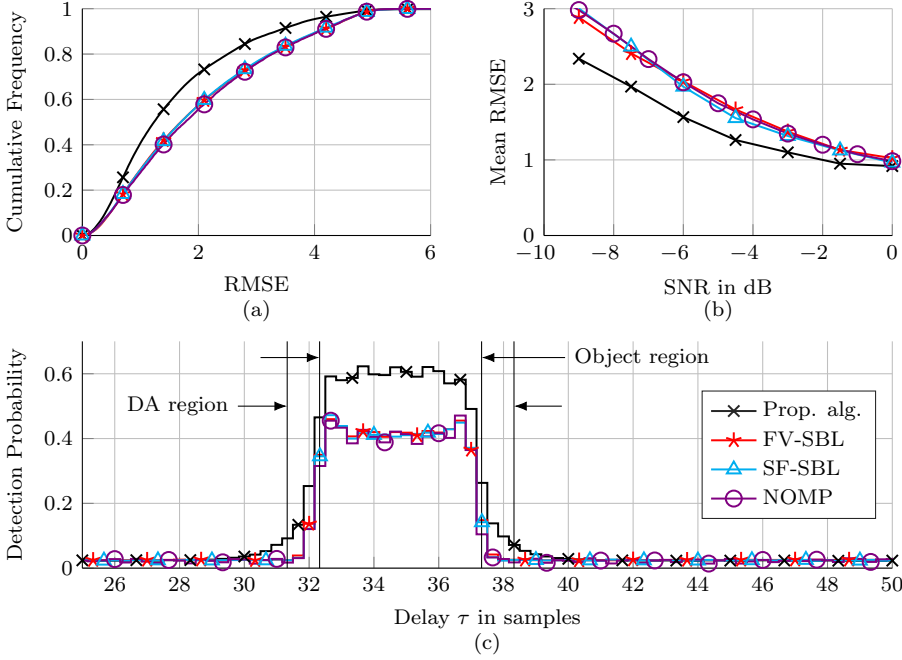


Figure E.4: Comparison of the proposed algorithm compared to the FV-SBL [45], SF-SBL [44] and NOMP [49]. Cumulative frequency of the RMSE at an SNR of -6 dB (a), mean RMSE over SNR (b) and histogram of the delay of the detected components at an SNR of -6 dB (c).

(E.33) with a few signal samples (“taps”) spaced with delay $\Delta\tau = \frac{1}{f_s}$. Thus, (E.34) is well approximated by

$$\mathbf{r} = \sum_{k=1}^K \sum_{l=0}^{L'_k} (\mathbf{s} \odot e^{-j2\pi(\tau_{\min,k} + l\Delta\tau)\mathbf{f}}) \alpha_{k,l} + \boldsymbol{\epsilon} \quad (\text{E.35})$$

where $\tau_{\min,k} = \min_l(\tau_{k,l})$ is the smallest delay of the k th target signal. Assuming that the pulse $s(t)$ has a flat spectrum, i.e. $\mathbf{s} = [1 \ 1 \cdots 1]^T$, the model (E.35) is an instance of (E.1). Thus, the proposed algorithm can be applied for the detection and estimation of the radar response from extended objects using $\theta_k = \tau_{\min,k}$ and $f(\theta_k, l) = \theta_k + l\Delta\tau$.

In order to showcase the ability of the algorithm to detect weak object signals and estimate their properties, we consider the case of a single object in noise. In this experiment, we perform 10^5 simulation runs. The object is modeled according to [7] with a uniform intensity function $q(\tau)$ such that on average 10 scatter points are drawn independently from the uniform PDF $U(\tau_{k,l}; 32.323T_s, 37.323)$. The amplitude of each scatter point are drawn independently from complex Gaussian PDFs $p(\alpha_{k,l}) = \text{CN}(\alpha_{k,l}; 0, 1)$ with zero mean and unit variance. In this setup, we select $\mathcal{S}_{\text{search}} = \{\min(\hat{S}_k) - 1, \min(\hat{S}_k), \dots, \max(\hat{S}_k) + 1\}$, set the thresholds to $\chi_1 = 3$ dB and $\chi_2 = 4$ dB. As comparison algorithms, we use the Newtonized

orthogonal matching pursuit (NOMP) algorithm¹⁴ [49], as well as two variants of SBL-based superresolution algorithms, the algorithm proposed in [45] abbreviated FV-SBL and the algorithm proposed in [44] abbreviated SF-SBL. Note that the SF-SBL is similar to the proposed algorithm using $\mathcal{S}_{\max} = \{0\}$ (i.e. allowing only a single component per group) and $\chi_1 = 0$ dB. To make the comparison fair, we selected the thresholds of the four algorithms such the mean number of components estimated outside the object region, was approximately the same. Numerical analysis revealed that this is achieved by thresholds of 6 dB for the FV-SBL and $\chi_2 = 3.5$ dB for the SF-SBL and a threshold (denoted as τ in [49]) of 6 dB for NOMP. For a fair comparison between the proposed algorithm and the comparison algorithms, we used an “oracle” data association (considered all components within ± 1 sample of the true object region to belong to the object) to estimate the extent \hat{O}_E and center-of-mass \hat{O}_C of the object. Note that the proposed algorithm inherently exploits its grouping mechanism increasing the detectability of components associated to an extended object. Let $\hat{\alpha}_{o,l}$ and $\hat{\tau}_{o,l}$ for $l \in \{1, 2, \dots, \hat{L}\}$ denote respective amplitudes and delays associated with the object, we estimated \hat{O}_E and \hat{O}_C as

$$\hat{O}_E = \max_l(\hat{\tau}_{o,l}) - \min_l(\hat{\tau}_{o,l}) \quad (\text{E.36})$$

$$\hat{O}_C = \frac{\sum_{l=1}^{\hat{L}} |\hat{\alpha}_{o,l}|^2 \hat{\tau}_{o,l}}{\sum_{l=1}^{\hat{L}} |\hat{\alpha}_{o,l}|^2}. \quad (\text{E.37})$$

Using (E.36) and (E.37), the root mean square error (RMSE) is given by $\text{RMSE} = \sqrt{(\hat{O}_E - O_E)^2 + (\hat{O}_C - O_C)^2}$ where O_E and O_C are the true object extent and center-of-mass. The cumulative frequency of the RMSE for an SNR of $\frac{\|r - \epsilon\|^2}{\|\epsilon\|^2} = -6$ dB is depicted in Figure E.4a while the mean RMSE over SNR is depicted in Figure E.4b. The proposed algorithm is able to outperform both SBL-based super-resolution methods in terms of RMSE at all SNR levels, although the difference is more pronounced in low-SNR conditions of -3 dB and less. The reason for this becomes evident when investigating the histogram of the locations of detected components depicted in Figure E.4c. While the mean number of components estimated in the object region is similar for the comparison methods, the proposed algorithm detects more components in the object region and is therefore able to estimate the parameters of the object more accurately. The increased number of estimated components is due to the lower threshold when adding new components within a group compared to adding new groups.

E.5.3 Variational Mode Decomposition

To demonstrate the flexibility of the proposed signal model and algorithm, we apply the proposed algorithm also to VMD [8]. VMD decomposes a signal \mathbf{x} into several “intrinsic mode functions” \mathbf{x}_k , i.e.,

$$\mathbf{x} = \sum_{k=1}^K \mathbf{x}_k. \quad (\text{E.38})$$

¹⁴<https://bitbucket.org/wcslspectralestimation/continuous-frequency-estimation/src/NOMP>

An “intrinsic mode function” $x_k(t) = A_k(t) \cos(\varphi_k(t))$ is defined in [8] as a sinusoidal function where the amplitude $A_k(t) \geq 0$ changes slowly over time t and the phase $\varphi_k(t)$ is a non-decreasing function with slowly time-varying instantaneous frequency $\frac{d\varphi_k(t)}{dt}$. The corresponding discrete signal model is given by $\mathbf{x}_k = [x_k(-\frac{N}{2}T_s) \ x_k((-\frac{N}{2} + 1)T_s) \ \cdots \ x_k((\frac{N}{2} - 1)T_s)]^T$, where \mathbf{x}_k is sampled with regular intervals $f_s = \frac{1}{T_s}$.

Note that the model (E.38) is not directly an instance of (E.1). The following differences are identified: (i) the signals \mathbf{x} and \mathbf{x}_k are real signals instead of complex ones, (ii) (E.39) does not consider the signal to be embedded in additive noise, and (iii) the signal model of VMD assumes the number of modes K is known whereas the signal model underlying the proposed algorithm assumes K to be unknown. The model difference in (i) can be bypassed by computing the discrete time analytic signal [56]. The model differences in (ii) and (iii) are considered a strength of the proposed model and algorithm, making it applicable to realistic and practical applications. It may not be immediately clear how to model each \mathbf{x}_k as a discrete sum of components, however, since $x_k(t)$ is essentially an amplitude and phase modulated signal, almost all of the signal energy is within a certain signal bandwidth that has to be much smaller than the sampling rate f_s . Therefore, each $x_k(t)$ can be well approximated by a signal of finite bandwidth, which (satisfying the sampling theorem) can be represented as a set of discrete frequencies spaced by $\Delta f = \frac{f_s}{N}$. In particular, we approximate the discrete-time analytic signal of each mode as $\mathbf{x}_{a,k} \approx \sum_{l=-L_k}^{L_k} \alpha_{k,l} e^{j2\pi f(\theta_k, l)t}$, where $f(\theta_k, l) = \theta_k + l\Delta f$ and L_k relates to the bandwidth B_k of the k th mode. The analytic signal \mathbf{x}_a is modeled as

$$\mathbf{x}_a \approx \sum_{k=1}^K \sum_{l=-L_k}^{L_k} \alpha_{k,l} e^{j2\pi f(\theta_k, l)t} + \boldsymbol{\epsilon} \quad (\text{E.39})$$

which is an instance of (E.1). The original signals (or estimates thereof) can be obtained as the real part of the corresponding analytical signal. We used a larger search space of $\mathcal{S}_{\text{search}} = \{\min(\hat{S}_k) - 5, \min(\hat{S}_k) - 4, \dots, \max(\hat{S}_k) + 5\}$ due to the larger signal length. We use a fixed threshold $\chi_1 = 3$ dB and varied the threshold $\chi_2 = 10$ dB and $\chi_2 = 18$ dB for the 10 dB SNR and 30 dB SNR case, respectively, based on preliminary investigations.

To demonstrate the estimation accuracy of the underlying modes, we generate a signal of length $N = 1000$ samples consisting of two modes in AWGN according to (E.38). The amplitude and instantaneous frequency are defined at support points $t_m \in \{-\frac{N}{2}T_s, -\frac{N}{4}T_s, 0, \frac{N}{4}T_s, (\frac{N}{2} - 1)T_s\}$ and linearly interpolated in between. The amplitude support points are defined as $A_{k,m} = A_{k,0}(1 + A_{\text{mod},k,m})$ where $A_{\text{mod},k,m}$ are random variables drawn independently for each mode k and support point m from a uniform PDF $U(A_{\text{mod},k,m}; -A_{\text{mod},k}, A_{\text{mod},k})$. Similarly, the instantaneous frequency at t_m is defined as $\frac{d\varphi_k(t)}{dt}|_{t=t_m} = 2\pi f_k(1 + f_{\text{mod},k,m})$ and linearly interpolated in between the points, where $f_{\text{mod},k,m}$ are i.i.d. uniform random variable with PDFs $U(f_{\text{mod},k,m}; -f_{\text{mod},k}, f_{\text{mod},k})$. The phase $\varphi_k[n]$ is obtained by integrating the instantaneous frequency from $t = -\frac{N}{2}T_s$ to nT_s . We select the modulation parameters for the first mode to be $f_k = 0.1$, $f_{\text{mod},k} = 0.66$, $A_{k,0} = 1$ and $A_{\text{mod},k} = 0.25$ and for the second mode we select $f_k = 0.35$, $f_{\text{mod},k} = 0.03$, $A_{k,0} = 1$ and $A_{\text{mod},k} = 0.9$. Thus, the amplitude modulation is more pronounced in the second mode and the frequency modulation is more pronounced in the first

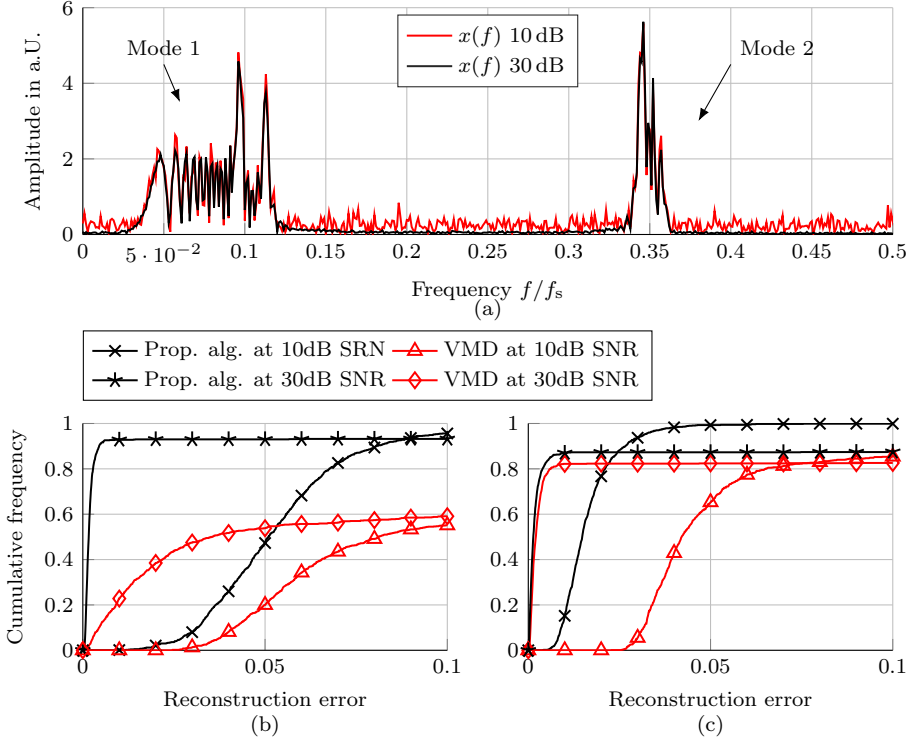


Figure E.5: Performance of the proposed algorithm for the task of VMD. (a) Example signal consisting of two modes, (b) cumulative frequency for the reconstruction error of mode 1 and (c) mode 2.

mode. An example signal resulting from these settings is shown in Figure E.5a.

We have generated (E.38) with added real-valued Gaussian noise ϵ of a SNR of 10 dB representing a noisy signal and a SNR of 30 dB representing a (nearly) noiseless signal and applied the proposed algorithm to the corresponding analytical signal and compared the mode estimation accuracy with the Matlab implementation of VMD (release R2022b). As performance metric we compared the estimation accuracy

$$\hat{e}_k = \min_l \frac{\|\mathbf{x}_k - \hat{\mathbf{x}}_l\|^2}{\|\mathbf{x}_k\|^2} \quad (\text{E.40})$$

where \mathbf{x}_k are the true modes of the signal \mathbf{x} and $\hat{\mathbf{x}}_l$ are the estimated modes. We performed 1000 simulation runs. Figures E.5b and E.5c depict the cumulative frequencies of the estimation error of both modes for all simulation realizations. The estimation performance of the proposed algorithm is better than the estimation performance of the VMD algorithm for both modes. However, it should be noted that the cumulative frequency of the estimation error “flattens out” at different values instead of converging towards 1, indicating some problem with the estimation process. This is due to the proposed algorithm overestimating the model order in 5 % and 19 % of the simulation runs for the 10 dB SNR case and 30 dB SNR

case, respectively. In these cases, an additional mode appears, resulting in a significantly decreased estimation performance. The VMD algorithm was designed for narrowband modes. Thus, one reason for the poor estimation performance of the VMD algorithm on Mode 1 could be the large bandwidth of this mode. However, the proposed method was able to outperform the VMD algorithm for both modes, i.e. also for the much narrower Mode 2.

Although we leave an extensive investigation of the performance of the proposed algorithm for VMD for future research, we chose to include these results to showcase how the algorithm can be applied to different problem settings.

E.6 Conclusion

In this paper, we introduced a generalization of SBL-based super-resolution methods for structured line spectra that is based on variational Bayesian inference. The key novelty of our proposed approach is to inherently model group-relation of certain spectral lines with a hierarchical Bernoulli-Gamma-Gaussian prior model. In the proposed probabilistic model, the occurrence of groups is regularized by a Bernoulli prior while the occurrence of spectral lines is regularized by a Jeffrey's prior on the amplitude precisions. Thus, a sparse estimate is obtained which automatically estimates the model order and the group structure in addition to the group parameters and component amplitudes.

The model can be straightforwardly adapted to a variety of inference problems based on the relation between the component frequencies $f(\theta_k, l)$ and the group parameter θ_k making it applicable to many diverse engineering problems. In particular, we demonstrate the superior performance of the proposed algorithm over other state-of-the-art algorithms based on three application examples, namely multi-pitch estimation, extended object detection based on radar signals, and variational mode decomposition (VMD). Most notably, in multi-pitch estimation the results based on the real Bach 10 dataset show that the proposed algorithm outperforms model-based algorithms and even pre-trained machine learning algorithms in terms of accuracy, precision, and recall. These three examples demonstrate the versatility of the proposed algorithm and show how incorporating knowledge about the structural relations between spectral lines into the estimation procedure can lead to performance gains.

In this work, we restrict the discussion to basis functions $\psi(\theta_k, l)$ of the form $e^{j2\pi f(\theta_k, l)t}$, i.e. spectral lines with frequency $f(\theta_k, l)$. However, as a promising direction for future research the underlying signal model can be extended many different set of basis functions making the proposed algorithm potentially applicable to an even wider variety of engineering problems. Another promising direction for future research is the extension of the proposed algorithm to sequential processing based on belief propagation message passing [63], e.g. for radio-frequency simultaneous localization and mapping [64, 65]. Another interesting avenue is using structured variational autoencoders [66] to learn arbitrary structured dictionaries.

E.7 Appendix

E.7.1 Update of q_λ

For the update of q_λ , we insert (E.7), (E.6) and $p(\lambda) = \text{Ga}(\lambda; \rho, \mu)$ into (E.10). Applying the logarithm on both sides, and ignoring all terms which do not depend on λ , we have

$$\ln p_\lambda \stackrel{e}{\propto} (N + \rho - 1) \ln \lambda - \lambda \langle \|\mathbf{x} - \hat{\Psi} \boldsymbol{\alpha}_{\hat{\mathcal{S}}}\|^2 \rangle_{q_\alpha} + \mu \quad (\text{E.41})$$

after a few algebraic manipulations. Solving the expectation of $\|\mathbf{x} - \hat{\Psi} \boldsymbol{\alpha}_{\hat{\mathcal{S}}}\|^2$ over q_α , we find that $q_\lambda = \text{Ga}(\lambda; N + \rho, \hat{M}_\lambda)$ is a Gamma distribution with shape $N + \rho$ and rate $\hat{M}_\lambda = \|\mathbf{x} - \hat{\Psi} \hat{\boldsymbol{\alpha}}\|^2 + \text{tr}(\hat{\Psi} \hat{\mathbf{C}} \hat{\Psi}^H) + \mu$, where $\text{tr}(\cdot)$ denote the matrix trace operator. The mean $\hat{\lambda}$ of this proxy PDF q_λ is

$$\hat{\lambda} = \frac{N + \rho}{\|\mathbf{x} - \hat{\Psi} \hat{\boldsymbol{\alpha}}\|^2 + \text{tr}(\hat{\Psi} \hat{\mathbf{C}} \hat{\Psi}^H) + \mu}. \quad (\text{E.42})$$

E.7.2 Variational Update of q_α

To find q_α we insert (E.12) into (E.13). After taking the logarithm on both sides and ignoring all terms which do not depend on α , we get

$$\ln q_\alpha(\alpha) \stackrel{e}{\propto} \langle \ln p(\mathbf{x} | \alpha, \hat{\boldsymbol{\theta}}, \lambda) + \ln p(\alpha | \gamma, \hat{\mathbf{z}}) \rangle_{q_\lambda q_\gamma}. \quad (\text{E.43})$$

Inserting (E.5) and (E.6) into (E.43) and ignoring all terms which do not depend on α , (E.43) can be rewritten as

$$\begin{aligned} \ln q_\alpha \stackrel{e}{\propto} & \langle -\lambda (\mathbf{x} - \hat{\Psi} \boldsymbol{\alpha}_{\hat{\mathcal{S}}})^H (\mathbf{x} - \hat{\Psi} \boldsymbol{\alpha}_{\hat{\mathcal{S}}}) - \boldsymbol{\alpha}_{\hat{\mathcal{S}}}^H \boldsymbol{\Gamma} \boldsymbol{\alpha}_{\hat{\mathcal{S}}} \\ & + \sum_{\alpha_{k,l} \notin \boldsymbol{\alpha}_{\hat{\mathcal{S}}}} \ln \delta(\alpha_{k,l}) \rangle_{q_\lambda q_\gamma}. \end{aligned} \quad (\text{E.44})$$

Using the expectation of q_λ , i.e., $\hat{\lambda} = \langle \lambda \rangle_{q_\lambda}$ (see (E.42)) and the mean of $q_{\gamma,k,l}$, i.e., $\hat{\gamma}_{k,l} = \langle \gamma_{k,l} \rangle_{q_{\gamma,k,l}}$ (see (E.23)) and denoting the real operator as $\text{Re}\{\cdot\}$, (E.44) can be rewritten as

$$\begin{aligned} \ln q_\alpha \stackrel{e}{\propto} & -(\boldsymbol{\alpha}_{\hat{\mathcal{S}}}^H (\hat{\lambda} \hat{\Psi}^H \hat{\Psi} + \hat{\boldsymbol{\Gamma}}) \boldsymbol{\alpha}_{\hat{\mathcal{S}}} - \text{Re}\{\cdot\} \boldsymbol{\alpha}_{\hat{\mathcal{S}}}^H 2\hat{\lambda} \hat{\Psi}^H \mathbf{x}) \\ & + \sum_{\alpha_{k,l} \notin \boldsymbol{\alpha}_{\hat{\mathcal{S}}}} \ln \delta(\alpha_{k,l}) \\ \stackrel{e}{\propto} & -(\boldsymbol{\alpha}_{\hat{\mathcal{S}}} - \hat{\boldsymbol{\alpha}})^H \hat{\mathbf{C}}^{-1} (\boldsymbol{\alpha}_{\hat{\mathcal{S}}} - \hat{\boldsymbol{\alpha}}) + \sum_{\alpha_{k,l} \notin \boldsymbol{\alpha}_{\hat{\mathcal{S}}}} \ln \delta(\alpha_{k,l}) \end{aligned} \quad (\text{E.45})$$

where we “complete the square” using $\hat{\mathbf{C}} = (\hat{\lambda} \hat{\Psi}^H \hat{\Psi} + \hat{\boldsymbol{\Gamma}})^{-1}$ and $\hat{\boldsymbol{\alpha}} = \hat{\lambda} \hat{\mathbf{C}} \hat{\Psi}^H \mathbf{x}$. Thus, we arrive at (E.15) and (E.16) with q_α being a complex Gaussian distribution of $\boldsymbol{\alpha}_{\hat{\mathcal{S}}}$ and $\alpha_{k,l} = 0$ for $\alpha_{k,l} \notin \boldsymbol{\alpha}_{\hat{\mathcal{S}}}$.

E.7.3 Derivation of $\ln Z(\theta, z)$ and $\Delta_k(\theta_k)$

We start by showing that expression I in (E.17) is the logarithm of a complex Gaussian distribution for $\alpha_{\bar{s}}$ and $\alpha_{k,l} = 0$ for $\alpha_{k,l} \notin \alpha_{\bar{s}}$ plus an expressions which depends on θ and z . Let $\Gamma_{\bar{s}} = \text{diag}(\gamma_{\bar{s}})$,

$$\begin{aligned} \text{I} &= \langle \ln p(\alpha, \theta, \gamma, z, \lambda | \mathbf{x}) \rangle_{q_{\lambda} q_{\gamma}} \\ &\stackrel{e}{\propto} \langle -\lambda(\mathbf{x} - \Psi_{\bar{s}} \alpha_{\bar{s}})^H (\mathbf{x} - \Psi_{\bar{s}} \alpha_{\bar{s}}) + \ln |\pi \Gamma_{\bar{s}}|^{-1} \\ &\quad - \alpha_{\bar{s}}^H \Gamma_{\bar{s}} \alpha_{\bar{s}} + \ln p(\theta) + \ln p(z) \\ &\quad + \sum_{\alpha_{k,l} \notin \alpha_{\bar{s}}} \ln \delta(\alpha_{k,l}) \rangle_{q_{\lambda} q_{\gamma}}. \end{aligned} \quad (\text{E.46})$$

After “completing the square” and adding $\ln |\pi C_{\bar{s}}|^{-1} - \ln |\pi C_{\bar{s}}|^{-1} = 0$ to complete the Gaussian distribution, we find

$$\begin{aligned} \text{I} &\stackrel{e}{\propto} \ln |\pi C_{\bar{s}}|^{-1} - (\alpha_{\bar{s}} - \hat{\alpha})^H C_{\bar{s}}^{-1} (\alpha_{\bar{s}} - \hat{\alpha}) \\ &\quad + \sum_{\alpha_{k,l} \notin \alpha_{\bar{s}}} \ln \delta(\alpha_{k,l}) + \hat{\lambda}^2 \mathbf{x}^H \Psi_{\bar{s}} C_{\bar{s}} \Psi_{\bar{s}}^H \mathbf{x} - \ln |\pi C_{\bar{s}}|^{-1} \\ &\quad + \langle \ln |\pi \Gamma_{\bar{s}}|^{-1} \rangle_{q_{\gamma}} + \ln p(\theta) + \ln p(z) \\ &\stackrel{e}{\propto} \ln \text{CN}(\alpha_{\bar{s}}; \hat{\alpha}, C_{\bar{s}}) + \sum_{\alpha_{k,l} \notin \alpha_{\bar{s}}} \ln \delta(\alpha_{k,l}) + \ln |C_{\bar{s}}| \\ &\quad + \hat{\lambda}^2 \mathbf{x}^H \Psi_{\bar{s}} C_{\bar{s}} \Psi_{\bar{s}}^H \mathbf{x} + \ln p(\theta) + \ln p(z) \\ &\quad + \sum_{\gamma_{k,l} \in \gamma_{\bar{s}}} \langle \ln \gamma_{k,l} \rangle_{q_{\gamma, k, l}}. \end{aligned} \quad (\text{E.47})$$

Inserting (E.47) into (E.17), all terms which depend on α integrate to 1 since they form a valid distribution. Thus, after integrating out α and taking the logarithm we arrive at (E.18).

To make the dependence of (E.18) on one set of parameters $\{\theta_k, z_k\}$ explicit, we express the covariance matrix $C_{\bar{s}}$ in $\ln Z([\hat{\theta}_{\bar{k}}, \theta_k], [\hat{z}_{\bar{k}}, 1])$ as a block matrix

$$C_{\bar{s}} = \begin{bmatrix} C_{\bar{k}}^{-1} & \hat{\lambda} \Psi_{\bar{k}}^H \Psi_k \\ \hat{\lambda} \Psi_k^H \Psi_{\bar{k}} & (\hat{\lambda} \Psi_k^H \Psi_k + \Gamma_k) \end{bmatrix}^{-1}$$

where $C_k = (\hat{\lambda} \Psi_k^H \Psi_k + \Gamma_k - \hat{\lambda}^2 \Psi_k^H \Psi_{\bar{k}} C_{\bar{k}} \Psi_{\bar{k}}^H \Psi_k)^{-1}$ and $C_{\bar{k}} = (\hat{\lambda} \Psi_{\bar{k}}^H \Psi_{\bar{k}} + \hat{\Gamma}_{\bar{k}})^{-1}$, to find $\ln |C_{\bar{s}}| = \ln |C_{\bar{k}}| + \ln |C_k|$ and, after applying the formula for block-matrix inversion,

$$\begin{aligned} &\hat{\lambda}^2 \mathbf{x}^H \Psi_{\bar{s}} C_{\bar{s}} \Psi_{\bar{s}}^H \mathbf{x} \\ &= \hat{\lambda}^4 \mathbf{x}^H \Psi_{\bar{k}} C_{\bar{k}} \Psi_{\bar{k}}^H \Psi_k C_k \Psi_k^H \Psi_{\bar{k}} C_{\bar{k}} \Psi_{\bar{k}}^H \mathbf{x} + \hat{\lambda}^2 \mathbf{x}^H \Psi_{\bar{k}} C_{\bar{k}} \Psi_{\bar{k}}^H \mathbf{x} \\ &\quad + \hat{\lambda}^2 \mathbf{x}^H \Psi_k C_k \Psi_k^H \mathbf{x} - \text{Re}\{\{\} 2 \hat{\lambda}^3 \mathbf{x}^H \Psi_k C_k \Psi_k^H \Psi_{\bar{k}} C_{\bar{k}} \Psi_{\bar{k}}^H \mathbf{x}\} \end{aligned}$$

which simplifies to

$$\hat{\lambda}^2 \mathbf{x}^H \Psi_{\bar{s}} C_{\bar{s}} \Psi_{\bar{s}}^H \mathbf{x} = \hat{\lambda}^2 \mathbf{x}^H \Psi_{\bar{k}} C_{\bar{k}} \Psi_{\bar{k}}^H \mathbf{x} + \mathbf{u}^H C_k^{-1} \mathbf{u}. \quad (\text{E.48})$$

Note, that $\mathbf{C}_{\hat{\mathcal{S}}}$ in the calculation of $\ln Z([\hat{\boldsymbol{\theta}}_{\bar{k}}, \theta_k], [\hat{\mathbf{z}}_{\bar{k}}, 0])$ equals $\mathbf{C}_{\bar{k}}$ in the calculation of $\ln Z([\hat{\boldsymbol{\theta}}_{\bar{k}}, \theta_k], [\hat{\mathbf{z}}_{\bar{k}}, 1])$. The same holds for $\boldsymbol{\Psi}_{\hat{\mathcal{S}}}$ and $\boldsymbol{\Psi}_{\bar{k}}$, respectively. Thus, we insert (E.48) into (E.18) and arrive at (E.19) after a few algebraic manipulations.

Bibliography

- [1] P. Stoica and R. Moses, *Spectral analysis of signals*. Upper Saddle River, NJ, USA: Pearson Prentice Hall, 2005.
- [2] E. Benetos, S. Dixon, Z. Duan, and S. Ewert, “Automatic music transcription: An overview,” *IEEE Signal Process. Mag.*, Jan. 2019, doi: 10.1109/MSP.2018.2869928.
- [3] M. Müller, D. P. W. Ellis, A. Klapuri, and G. Richard, “Signal processing for music analysis,” *IEEE J. Sel. Topics Signal Process.*, Feb. 2011, doi: 10.1109/JSTSP.2011.2112333.
- [4] M. G. Christensen and A. Jakobsson, *Multi-Pitch Estimation*, ser. Synthesis Lectures on Speech & Audio Processing, B. H. Juang, Ed. San Rafael, CA, USA: Morgan & Claypool, 2009.
- [5] M. G. Christensen, P. Stoica, A. Jakobsson, and S. Holdt Jensen, “Multi-pitch estimation,” *Signal Process.*, Apr. 2008, doi: <https://doi.org/10.1016/j.sigpro.2007.10.014>.
- [6] K. Granstrom, M. Baum, and S. Reuter, “Extended object tracking: Introduction, overview, and applications,” *J. Advances Inf. Fusion*, vol. 12, no. 2, pp. 139–174, Dec. 2017.
- [7] F. M. Schubert, M. L. Jakobsen, and B. H. Fleury, “Non-stationary propagation model for scattering volumes with an application to the rural LMS channel,” *IEEE Trans. Antennas Propag.*, Jan. 2013, doi: 10.1109/TAP.2013.2242821.
- [8] K. Dragomiretskiy and D. Zosso, “Variational mode decomposition,” *IEEE Trans. Signal Process.*, Feb. 2014, doi: 10.1109/TSP.2013.2288675.
- [9] R. Schmidt, “Multiple emitter location and signal parameter estimation,” *IEEE Trans. Antennas Propag.*, Mar. 1986, doi: 10.1109/TAP.1986.1143830.
- [10] R. Roy and T. Kailath, “ESPRIT-estimation of signal parameters via rotational invariance techniques,” *IEEE Trans. Acoust., Speech, Signal Process.*, Jul. 1989, doi: 10.1109/29.32276.
- [11] I. Ziskind and M. Wax, “Maximum likelihood localization of multiple sources by alternating projection,” *IEEE Trans. Acoust., Speech, Signal Process.*, Oct. 1988, doi: 10.1109/29.7543.
- [12] M. Feder and E. Weinstein, “Parameter estimation of superimposed signals using the EM algorithm,” *IEEE Trans. Acoust., Speech, Signal Process.*, Apr. 1988, doi: 10.1109/29.1552.

- [13] P. Stoica and Y. Selen, "Model-order selection: a review of information criterion rules," *IEEE Signal Process. Mag.*, Jul. 2004, doi: 10.1109/MSP.2004.1311138.
- [14] R. Tibshirani, "Regression shrinkage and selection via the LASSO," *J. Roy. Statistical Soc.: Ser. B (Statistical Methodology)*, 1996, doi: <https://doi.org/10.1111/j.2517-6161.1996.tb02080.x>.
- [15] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Rev.*, Mar. 2001, doi: 10.1137/S003614450037906X.
- [16] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. Signal Process.*, 1993, doi: 10.1109/78.258082.
- [17] M. Tipping, "The relevance vector machine," in *Advances Neural Inf. Process. Syst.*, vol. 12. Denver, CO, USA: MIT Press, Nov. 29 – Dec. 4, 1999, pp. 652–658.
- [18] M. E. Tipping and A. C. Faul, "Fast marginal likelihood maximisation for sparse Bayesian models," in *Proc. 9th Int. Workshop Artif. Intell. and Statist.*, vol. R4, Key West, FL, USA, Jan. 03–06, 2003, pp. 276–283.
- [19] D. Shutin, T. Buchgraber, S. R. Kulkarni, and H. V. Poor, "Fast variational sparse Bayesian learning with automatic relevance determination for superimposed signals," *IEEE Trans. Signal Process.*, Dec. 2011, doi: 10.1109/TSP.2011.2168217.
- [20] P. Stoica, P. Babu, and J. Li, "SPICE: A sparse covariance-based estimation method for array processing," *IEEE Trans. Signal Process.*, Feb. 2011, doi: 10.1109/TSP.2010.2090525.
- [21] M. Yuan and Y. Lin, "Model selection and estimation in regression with grouped variables," *J. Roy. Statistical Soc.: Ser. B (Statistical Methodology)*, Feb. 2006, doi: <https://doi.org/10.1111/j.1467-9868.2005.00532.x>.
- [22] M. Kyung, J. Gill, M. Ghosh, and G. Casella, "Penalized regression, standard errors, and Bayesian LASSOs," *Bayesian Anal.*, Jun. 2010, doi: 10.1214/10-BA607.
- [23] S. Raman, T. J. Fuchs, P. J. Wild, E. Dahl, and V. Roth, "The Bayesian group-LASSO for analyzing contingency tables," in *Proc. 26th Annu. Int. Conf. Mach. Learn.*, Jun. 14–18, 2009, doi: 10.1145/1553374.1553487.
- [24] X. Xu and M. Ghosh, "Bayesian variable selection and estimation for group LASSO," *Bayesian Anal.*, Dec. 2015, doi: 10.1214/14-BA929.
- [25] Y. Kim, J. Kim, and Y. Kim, "Blockwise sparse regression," *Statistica Sinica*, vol. 16, no. 2, pp. 375–390, Apr. 2006.
- [26] Y. C. Eldar, P. Kuppinger, and H. Bolcskei, "Block-sparse signals: Uncertainty relations and efficient recovery," *IEEE Trans. Signal Process.*, Jun. 2010, doi: 10.1109/TSP.2010.2044837.

- [27] Z. Zhang and B. D. Rao, "Sparse signal recovery with temporally correlated source vectors using sparse Bayesian learning," *IEEE J. Sel. Topics Signal Process.*, 2011, doi: 10.1109/JSTSP.2011.2159773.
- [28] —, "Extension of SBL algorithms for the recovery of block sparse signals with intra-block correlation," *IEEE Trans. Signal Process.*, Apr. 2013, doi: 10.1109/TSP.2013.2241055.
- [29] Z. Ma, W. Dai, Y. Liu, and X. Wang, "Group sparse Bayesian learning via exact and fast marginal likelihood maximization," *IEEE Trans. Signal Process.*, May 2017, doi: 10.1109/TSP.2017.2675867.
- [30] D. P. Wipf and B. D. Rao, "An empirical Bayesian strategy for solving the simultaneous sparse approximation problem," *IEEE Trans. Signal Process.*, Jun. 2007, doi: 10.1109/TSP.2007.894265.
- [31] S. D. Babacan, S. Nakajima, and M. N. Do, "Bayesian group-sparse modeling and variational inference," *IEEE Trans. Signal Process.*, Jun. 2014, doi: 10.1109/TSP.2014.2319775.
- [32] J. Möderl, F. Pernkopf, K. Witrisal, and E. Leitinger, "Fast variational block-sparse Bayesian learning," *ArXiv e-prints*, Jun. 2023, doi: 10.48550/arXiv.2306.00442.
- [33] J. Fang, Y. Shen, H. Li, and P. Wang, "Pattern-coupled sparse Bayesian learning for recovery of block-sparse signals," *IEEE Trans. Signal Process.*, Jan. 2015, doi: 10.1109/TSP.2014.2375133.
- [34] T. Kronvall, S. I. Adalbjörnsson, S. Nadig, and A. Jakobsson, "Group-sparse regression using the covariance fitting criterion," *Signal Process.*, Oct. 2017, doi: <https://doi.org/10.1016/j.sigpro.2017.03.025>.
- [35] D. P. Wipf and B. D. Rao, "Sparse Bayesian learning for basis selection," *IEEE Trans. Signal Process.*, Aug. 2004, doi: 10.1109/TSP.2004.831016.
- [36] D. P. Wipf, B. D. Rao, and S. Nagarajan, "Latent variable Bayesian models for promoting sparsity," *IEEE Trans. Inf. Theory*, Sep. 2011, doi: 10.1109/TIT.2011.2162174.
- [37] D. Wipf, "Sparse estimation with structured dictionaries," in *Advances Neural Inf. Process. Syst.*, vol. 24. Granada, Spain: Curran Associates, Dec. 2011, pp. 2016–2024.
- [38] Y. Chi, L. L. Scharf, A. Pezeshki, and A. R. Calderbank, "Sensitivity to basis mismatch in compressed sensing," *IEEE Trans. Signal Process.*, May 2011, doi: 10.1109/TSP.2011.2112650.
- [39] M. F. Duarte and R. G. Baraniuk, "Spectral compressive sensing," *Appl. Comput. Harmon. Anal.*, Jul. 2013, doi: <https://doi.org/10.1016/j.acha.2012.08.003>.
- [40] M.-A. Badiu, T. L. Hansen, and B. H. Fleury, "Variational Bayesian inference of line spectra," *IEEE Trans. Signal Process.*, May 2017, doi: 10.1109/TSP.2017.2655489.

- [41] Z. Yang and L. Xie, “On gridless sparse methods for line spectral estimation from complete and incomplete data,” *IEEE Trans. Signal Process.*, Jun. 2015, doi: 10.1109/TSP.2015.2420541.
- [42] S. Grebien, E. Leitinger, B. H. Fleury, and K. Witrisal, “Super-resolution channel estimation including the dense multipath component — A sparse variational Bayesian approach,” *ArXiv e-prints*, 2023. [Online]. Available: <https://arxiv.org/abs/2308.01702>
- [43] T. L. Hansen, M. A. Badiu, B. H. Fleury, and B. D. Rao, “A sparse Bayesian learning algorithm with dictionary parameter estimation,” in *2014 IEEE 8th Sensor Array and Multichannel Signal Process. Workshop (SAM)*, Jun. 22–25, 2014, doi: 10.1109/SAM.2014.6882422.
- [44] T. L. Hansen, B. H. Fleury, and B. D. Rao, “Superfast line spectral estimation,” *IEEE Trans. Signal Process.*, Feb. 2018, doi: 10.1109/TSP.2018.2807417.
- [45] D. Shutin, W. Wand, and T. Jost, “Incremental sparse Bayesian learning for parameter estimation of superimposed signals,” in *10th Int. Conf. Sampling Theory and Appl.*, Bremen, Germany, Jul. 1–5, 2013, pp. 513–516.
- [46] S. I. Adalbjörnsson, A. Jakobsson, and M. G. Christensen, “Multi-pitch estimation exploiting block sparsity,” *Signal Process.*, Apr. 2015, doi: <https://doi.org/10.1016/j.sigpro.2014.10.014>.
- [47] J. Swärd, H. Li, and A. Jakobsson, “Off-grid fundamental frequency estimation,” *IEEE/ACM Trans. Audio, Speech, Language Process.*, Feb. 2018, doi: 10.1109/TASLP.2017.2775800.
- [48] 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.*, Sep. 2016, doi: 10.1109/TSP.2016.2572041.
- [49] B. Mamandipoor, D. Ramasamy, and U. Madhow, “Newtonized orthogonal matching pursuit: Frequency estimation over the continuum,” *IEEE Trans. Signal Process.*, Oct. 2016, doi: 10.1109/TSP.2016.2580523.
- [50] Z. Duan and B. Pardo. Bach 10 dataset. Accessed: Dec. 2022. [Online]. Available: www.ece.rochester.edu/projects/air/datasets/Bach10.zip
- [51] E. Benetos and T. Weyde, “An efficient temporally-constrained probabilistic model for multiple-instrument music transcription,” in *Proc. 16th Int. Soc. Music Inf. Retrieval Conf.*, Malaga, Spain, Oct. 26–30, 2015.
- [52] D. G. Tzikas, A. C. Likas, and N. P. Galatsanos, “The variational approximation for Bayesian inference,” *IEEE Signal Process. Mag.*, Nov. 2008, doi: 10.1109/MSP.2008.929620.
- [53] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2006.

- [54] M. Titsias and M. Lázaro-Gredilla, “Spike and slab variational inference for multi-task and multiple kernel learning,” in *Advances Neural Inf. Process. Syst.*, vol. 24. Granada, Spain: Curran Associates, Dec. 2011, p. 2339. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2011/file/b495ce63ede0f4efc9eec62cb947c162-Paper.pdf
- [55] E. Leitinger, S. Grebien, B. Fleury, and K. Witrisal, “Detection and estimation of a spectral line in MIMO systems,” in *2020 54th Asilomar Conf. Signals, Syst. and Computers*, Nov. 01–04, 2020, doi: 10.1109/IEEECONF51394.2020.9443381.
- [56] L. Marple, “Computing the discrete-time “analytic” signal via FFT,” *IEEE Trans. Signal Process.*, Sep. 1999, doi: 10.1109/78.782222.
- [57] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, “A consistent metric for performance evaluation of multi-object filters,” *IEEE Trans. Signal Process.*, Aug. 2008, doi: 10.1109/TSP.2008.920469.
- [58] Z. Duan, B. Pardo, and C. Zhang, “Multiple fundamental frequency estimation by modeling spectral peaks and non-peak regions,” *IEEE/ACM Trans. Audio, Speech, Language Process.*, Nov. 2010, doi: 10.1109/TASL.2010.2042119.
- [59] A. de Cheveigné and H. Kawahara, “YIN, a fundamental frequency estimator for speech and music,” *J. Acoust. Soc. Amer.*, Apr. 2022, doi: <https://doi.org/10.1121/1.458024>.
- [60] F. Elvander, J. Swärd, and A. Jakobsson, “Online estimation of multiple harmonic signals,” *IEEE/ACM Trans. Audio, Speech, Language Process.*, Feb. 2017, doi: 10.1109/TASLP.2016.2634118.
- [61] M. Bay, A. F. Ehmann, and J. S. Downie, “Evaluation of multiple-F0 estimation and tracking systems,” in *Proc. 10th Int. Soc. Music Inf. Retrieval Conf.*, Kobe, Japan, Oct. 26 – 30, 2009, pp. 315–320.
- [62] X. Li, Y. Yan, J. Soraghan, Z. Wang, and J. Ren, “A music cognition-guided framework for multi-pitch estimation,” *Cogn. Comput.*, pp. 1–13, 2022.
- [63] X. Li, E. Leitinger, A. Venus, and F. Tufvesson, “Sequential detection and estimation of multipath channel parameters using belief propagation,” *IEEE Trans. Wireless Commun.*, Oct. 2022, doi: 10.1109/TWC.2022.3165856.
- [64] E. Leitinger, F. Meyer, F. Hlawatsch, K. Witrisal, F. Tufvesson, and M. Z. Win, “A belief propagation algorithm for multipath-based SLAM,” *IEEE Trans. Wireless Commun.*, Dec. 2019, doi: 10.1109/TWC.2019.2937781.
- [65] E. Leitinger, A. Venus, B. Teague, and F. Meyer, “Data fusion for multipath-based SLAM: Combining information from multiple propagation paths,” *IEEE Trans. Signal Process.*, Sep. 2023, doi: 10.1109/TSP.2023.3310360.
- [66] M. J. Johnson, D. K. Duvenaud, A. Wiltchko, R. P. Adams, and S. R. Datta, “Composing graphical models with neural networks for structured representations and fast inference,” in *Advances Neural Inf. Process. Syst.*, vol. 29. Barcelona, Spain: Curran Associates, Inc., Dec. 5–10, 2016.

