



HAL
open science

Approximations variationnelles bayésiennes pour le traitement du signal

Angélique Drémeau

► **To cite this version:**

Angélique Drémeau. Approximations variationnelles bayésiennes pour le traitement du signal: quelques applications en parcimonie et reconstruction de phase. Mathématiques [math]. UBO - Université de Brest, 2022. tel-04454630

HAL Id: tel-04454630

<https://ensta-bretagne.hal.science/tel-04454630>

Submitted on 13 Feb 2024

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

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

HABILITATION À DIRIGER LES RECHERCHES DE

L'UNIVERSITÉ DE BRETAGNE OCCIDENTALE

ÉCOLE DOCTORALE N° 601
*Mathématiques et Sciences et Technologies
de l'Information et de la Communication*

Par

Angélique DRÉMEAU

Approximations variationnelles bayésiennes pour le traitement du signal :

quelques applications en parcimonie et reconstruction de phase

HDR présentée et soutenue à l'ENSTA Bretagne, le 1er décembre 2022
Unité de recherche : Lab-STICC, UMR CNRS 6285

Rapporteurs avant soutenance :

Christian Jutten	Professeur des Universités, Université Grenoble Alpes, Laboratoire Gipsa-Lab UMR CNRS 5216
Pascal Larzabal	Professeur des Universités, IUT de Cachan, Laboratoire SATIE UMR CNRS 8069
Jérôme Mars	Professeur des Universités, Université Grenoble-INP, Laboratoire Gipsa-Lab UMR CNRS 5216

Composition du Jury :

Président :	Emanuel Radoi	Fonction et établissement d'exercice (<i>à préciser après la soutenance</i>)
Examineurs :	Emanuel Radoi	Professeur des Universités, UBO, Lab-STICC UMR CNRS 6285
	Isabelle Quidu	Maître de conférence, HDR, UBO, Lab-STICC UMR CNRS 6285

Invité(s) :

Laurent Daudet	Professeur des Universités, Université Paris Cité, Institut Langevin UMR CNRS 7587
Gaultier Real	Expert technique environnement et performances LSM, DGA Techniques Navales

À Ève,
projet de recherche ambitieux

Résumé

Alternatives aux méthodes d'échantillonnage, les approximations variationnelles bayésiennes constituent une famille d'approximations approchant des distributions de probabilité par d'autres, présentant des factorisations désirables. Mises en oeuvre dans des procédures algorithmiques peu complexes, elles permettent d'obtenir une expression analytique (quoique approchée) des distributions cibles. Ces méthodes sont considérées avec succès pour la résolution de problèmes inverses dans des cadres bayésiens.

Fil conducteur de mes recherches, les approximations variationnelles bayésiennes ont notamment prouvé leur intérêt dans deux grands types de problèmes :

- les décompositions parcimonieuses (structurées),
- la reconstruction de phase (informée).

Mes recherches m'ont permis d'aborder ces deux grands problèmes selon deux points de vue : un point de vue méthodologique, s'intéressant au développement de nouvelles solutions basées sur des approximations variationnelles bayésiennes, et un point de vue appliqué, mettant en oeuvre ces solutions dans des cadres pratiques concrets. Parmi ces applications, l'acoustique sous-marine tient une place importante, à travers notamment le traitement d'antenne en milieu océanique fluctuant ou l'estimation modale en ultra-basses fréquences et milieux petit-fond.

Ces travaux ouvrent la voie à de nombreuses perspectives de recherche tant méthodologiques qu'applicatives, combinant modélisation bayésienne, apprentissage automatique et prise en compte de la physique des phénomènes considérés.

Abstract

As an alternative to sampling methods, Bayesian variational approximations are a family of approximations that approximate probability distributions by others, with desirable factorizations. Implemented in low complexity algorithmic procedures, they allow to obtain an analytical expression (although approximated) of the target distributions. These methods are successfully considered for solving inverse problems in Bayesian frameworks.

At the heart of my research, Bayesian variational approximations have notably proven their value in two main types of problems :

- (structured) sparse decompositions,
- (informed) phase retrieval.

My research has allowed me to approach these two major problems from two points of view : a methodological point of view, focusing on the development of new solutions based on Bayesian variational approximations, and an applied point of view, implementing these solutions in concrete practical settings. Among these applications, underwater acoustics holds an important place, notably through antenna processing in fluctuating oceanic environments or modal estimation in ultra-low frequencies and shallow-water environments.

This work opens the way to numerous research perspectives, both methodological and applicative, combining Bayesian modeling, machine learning and taking into account the physics of the phenomena considered.

Avant-propos : du plaisir dans la recherche

En matière d'habilitation à diriger des recherches, il existe autant de modèles de manuscrits qu'il y a de manuscrits. Pour rédiger le mien – autant être honnête – je ne me suis donné qu'une seule ligne de conduite : me faire plaisir. Mais quelque part, n'est-ce pas le propre du (de la) chercheur(se) d'agir uniquement pour son propre plaisir ? Un(e) chercheur(se) n'est-il(elle) pas l'être le plus égoïste ? Ne place-t'il(elle) pas son plaisir personnel avant toute idée d'utilité ?

Je l'avoue sans scrupule. J'ai choisi la recherche pour ce qu'elle offre d'apprentissage, d'enrichissement intellectuel,... un statut d'étudiante à vie. Et, entre nous, mieux vaut trouver en soi les motivations de ce métier. Quelles autres en effet ? Repousser les limites de l'ignorance humaine ? Voilà un objectif bien ambitieux, bien prétentieux qui peut mener à de belles déconvenues ! Attendre de la reconnaissance, chercher à atteindre les hautes strates des distinctions ? En d'autres temps, peut-être, mais voyez comme la parole scientifique est sans cesse remise en question dans l'espace public, ramenée à un avis parmi d'autres, sans respect, sans reconnaissance d'une expertise quelconque, sans égard ne serait-ce que pour le temps passé à l'élaboration du raisonnement... Non, j'en suis persuadée, la voie du plaisir est la seule tenable : le chercheur en est son seul garant, il ne tient qu'à lui d'être heureux...

Mais si sa démarche est effectivement très égoïste, le chercheur n'est pour autant pas isolé. A côté de lui, de cet éternel étudiant donc, il y a d'autres étudiants... plus jeunes. Ceux qu'il encadre, qu'il suit, de la licence ou du diplôme d'ingénieur à la thèse. Ne vous y trompez pas : un plaisir égoïste là encore. Plus immédiat, plus franc, plus objectif, moins trompeur. J'ai la chance d'exercer le métier d'enseignante en plus de celui de chercheuse. Sept ans se sont écoulés depuis mon premier amphi, et c'est toujours avec la boule au ventre, la gorge serrée que je me présente devant la promotion à chacun de mes cours magistraux. Masochiste, l'enseignante-chercheuse ? Oui, sans doute un peu aussi. Mais quelles ailes me poussent lorsqu'un étudiant – un seul oui, cela suffit – vient me trouver, prend quelques minutes de son temps, pour me dire qu'il a apprécié mon cours ! Je deviens alors fière, invincible, prête à affronter n'importe quel formulaire de dépôt de projet.

Dans ce manuscrit, j'ai tenté de ne pas trop empiéter sur le travail de mes étudiants – les plus grands, ceux que j'ai encadrés en thèse. Car si ce manuscrit a pour objectif de valider ma capacité à diriger des recherches, il est important, je crois, de rappeler que pour diriger, il faut savoir où on va et comment on y va, et pas seulement pointer le doigt en disant "c'est par là". Dans mon idée de me "faire plaisir", j'ai donc pris ce manuscrit comme prétexte pour faire un focus sur ce que je pense être un fil rouge de mes recherches : les approximations variationnelles bayésiennes. Ces recherches, je les ai d'abord commencées en collaboration avec Cédric Herzet et Laurent Daudet, alors que j'étais en post-doctorat à l'Institut Langevin. Je les ai poursuivies avec "mes" doctorants aujourd'hui docteurs, Guillaume Beaumont, qui a essuyé les plâtres de mon encadrement balbutiant¹, Thomas Paviet-

1. Merci à Ronan Fablet d'avoir si bien endossé le rôle de directeur de thèse rassurant tout le monde.

Salomon, qui a eu le courage de me suivre en thèse après m'avoir subie en cours², et avec Clément Dorffer, post-doctorant lisant à la fois dans la matrice et dans mes idées. Avec eux, j'ai plongé dans l'univers de l'acoustique sous-marine et découvert des ramifications qui m'ont conduite à quitter le domaine du bayésien pour celui de la physique ou de la bioacoustique. Après eux, ont ainsi démarré les thèses de Fabio Cassiano et Alexandre L'Her qui m'emmènent vers des terres moins familières mais tout aussi riches, et, depuis novembre dernier, celle de Perrine Bauchot avec qui j'accepte enfin de succomber aux sirènes du machine learning.

L'habilitation à diriger des recherches arrive à point nommé. Il me faut prendre du recul, lever la tête du guidon, écrire un nouveau projet de recherche. J'ai eu la chance de pouvoir mener à bien déjà celui que j'ai écrit pour mon intégration à l'ENSTA Bretagne en janvier 2015. Il me semble primordial ici de remercier mon institution, justement, qui a fait montre d'un soutien sans faille dans le montage des projets qui me tenaient à cœur. Depuis mon arrivée à l'ENSTA Bretagne, j'ai également été soutenue par DGA Techniques Navales, il me faut remercier Dominique Fattaccioli et Gaultier Real d'avoir eu confiance dans mes pistes de recherche.

Un nouveau projet de recherche exige de voir plus loin que les simples avancées incrémentales qui s'imposent inévitablement à l'issue d'un projet, d'une thèse. Pour un chercheur, cela peut donner le tournis et je me suis imaginée un instant bifurquer dans un domaine totalement nouveau... Mais, une fois n'est pas coutume, je vais mettre de côté ma tendance naturelle au plaisir égoïste et tenter de m'appuyer sur ce que je sais déjà pour faire modestement avancer la "science".
Au moins un peu.

2. Merci à Thierry Chonavel d'avoir accepté d'être directeur de thèse dans un encadrement multiple!

Table des matières

Résumé	5
Abstract	7
Avant-propos	9
1 Introduction	13
1.1 Approximations variationnelles bayésiennes	13
1.2 Contributions et plan du document	14
1.3 Notations	15
2 Décompositions parcimonieuses	17
2.1 Contribution méthodologique	17
2.1.1 Algorithmes gloutons revisités (travaux de thèse)	18
2.1.2 Algorithmes variationnels bayésiens	19
2.1.3 Machine de Boltzmann	19
2.1.4 Article sélectionné	20
2.2 Application : estimation modale en ASM UBF	21
2.2.1 SoBaP informé	22
2.2.2 Travaux ultérieurs et concomitants	23
2.2.3 Articles sélectionnés	24
2.3 Synthèse et ouverture	24
3 Reconstruction de phase	27
3.1 Contribution méthodologique	27
3.2 Reconstruction aveugle : propagation d'ondes optiques en milieux complexes	29
3.2.1 Calibration du milieu diffusant	29
3.2.2 Focalisation à travers le milieu diffusant	30
3.2.3 Article sélectionné	32
3.3 Reconstruction informée : estimation de DOA en présence de bruit de phase	32
3.3.1 Travaux ultérieurs et concomittants	32
3.3.2 Article sélectionné	33
3.4 Synthèse et ouverture	33
4 Perspectives de recherche	35
4.1 De ce que je n'ai pas écrit jusque là	35
4.2 Projet de recherche	36
4.2.1 Modèles <i>a priori</i> structurés	36

4.2.2	Modèles d'observation non-linéaires	38
4.3	Synthèse et ouverture	40
	Bibliographie	41
	Annexes	44
	A Curriculum Vitae	45
	B Publications et communications	51
	C Articles sélectionnés	57
	– Boltzmann machine and mean-field approximation for structured sparse decompositions	58
	– Reconstruction of dispersion curves in the frequency-wavenumber domain using compressed sensing on a random array	72
	– Modal estimation in underwater acoustics by data-driven structured sparse decompositions	81
	– Reference-less measurement of the transmission matrix of a highly scattering material using a DMD and phase retrieval techniques	86
	– DOA estimation in structured phase-noisy environments	100

Chapitre 1

Introduction

Mes travaux de recherche ont pour thématique générale la résolution de problèmes inverses, les plus souvent posés dans un cadre bayésien. Ma thèse, mes post-doctorats et mes recherches actuelles portent, dans leur grande majorité, sur différentes instances de cette large thématique.

Ainsi, je me suis intéressée pendant ma thèse à la compression d'image, puis, pendant mes quatre années de post-doctorats, à la représentation de sons, l'analyse de données multimédia et l'imagerie optique. Depuis janvier 2015 et mon intégration à l'ENSTA Bretagne, je travaille sur des problèmes de traitement du signal acoustique sous-marin, plus particulièrement sur la localisation de sources et l'estimation de modes propagatifs en ultra-basses fréquences. Etant pour la plupart mal posés, ces problèmes m'ont amenée à introduire différents types d'information *a priori* – des modèles probabilistes encourageant la parcimonie (comme les modèles Bernoulli-gaussiens ou les machines de Boltzmann) à des modèles probabilistes formalisant des dépendances temporelles (comme les processus gaussiens ou les chaînes de Markov). Enfin, trait d'union entre mon dernier post-doctorat et mes recherches actuelles – et sujet qui me tient particulièrement à coeur –, je suis confrontée à des problèmes interrogeant les modèles d'observation, dépassant le cadre largement répandu du modèle linéaire.

Dans ce contexte, mes contributions s'attachent à la conception d'algorithmes de résolution performants et peu coûteux en complexité algorithmique. Exploitant les modèles parcimonieux, j'ai par exemple proposé de nouveaux algorithmes de décomposition, d'apprentissage de dictionnaire ou d'inking spatial.

Fil conducteur de mes travaux, le cadre bayésien m'a permis de recourir à des outils probabilistes efficaces. A l'issue de ma thèse, je me suis plus particulièrement penchée sur les approximations variationnelles bayésiennes pour l'estimation de distributions *a posteriori*. Ces techniques éprouvées m'ont permis de développer des solutions à la fois performantes et raisonnables en terme de coût calculatoire.

1.1 Approximations variationnelles bayésiennes

Les approximations variationnelles bayésiennes font référence à une famille d'approximations approchant des distributions de probabilité *a posteriori* par des distributions présentant des factorisations désirables. Le terme "désirable" peut être entendu de façons différentes; en général, on souhaite que ces factorisations autorisent des formes analytiques et/ou des procédures de calcul simples.

Formellement, ces approximations sont calculées comme la solution de problèmes de minimisation mettant en jeu la divergence de Kullback-Leibler avec la distribution *a posteriori* cible [42]. Notant $\mathbf{u} \triangleq [u_1, \dots, u_I]^T$ la variable (de dimension I) de dépendance d'observations notées \mathbf{y} , l'approximation

$q(\mathbf{u})^*$ de la loi *a posteriori* $p(\mathbf{u}|\mathbf{y})$ sera ainsi définie comme

$$q(\mathbf{u})^* = \arg \min_{q \in \mathcal{Q}} \int_{\mathbf{u}} q(\mathbf{u}) \log \frac{q(\mathbf{u})}{p(\mathbf{u}|\mathbf{y})} d\mathbf{u}, \quad (1.1)$$

où \mathcal{Q} représente un ensemble de contraintes sur l'espace des densités de probabilité. Selon cet ensemble \mathcal{Q} , la minimisation (1.1) donne lieu à des approximations différentes.

Dans mes travaux et ceux que j'ai pu encadrer, on s'intéresse à deux approximations particulières :

- L'approximation de champ moyen fait l'hypothèse d'une indépendance *a posteriori* de tous les éléments de \mathbf{u} :

$$\mathcal{Q}_{\text{MF}} = \left\{ q \mid q(\mathbf{u}) = \prod_{i=1}^I q_i(u_i) \right\}. \quad (1.2)$$

- L'approximation de Bethe exploite l'hypothèse que la distribution jointe $p(\mathbf{u}, \mathbf{y})$ admet une factorisation reposant sur une partition en E régions des variables d'intérêt $\{u_i\}_{i=1}^I$, soit

$$p(\mathbf{u}, \mathbf{y}) = \frac{1}{Z} \prod_{\epsilon=1}^E f_{\epsilon}(\mathbf{u}_{\epsilon}, \mathbf{y}), \quad (1.3)$$

où Z est une constante de normalisation et les facteurs $f_{\epsilon}(\mathbf{u}_{\epsilon}, \mathbf{y})$ dépendent des \mathbf{u}_{ϵ} , sous-ensembles de $\{u_i\}_{i=1}^I$ et des observations \mathbf{y} . Sous cette hypothèse, l'approximation de Bethe propose de définir l'espace des contraintes $\mathcal{Q}_{\text{Bethe}}$ comme

$$\mathcal{Q}_{\text{Bethe}} = \left\{ q \mid q(\mathbf{u}) = \frac{\prod_{\epsilon=1}^E q_{\epsilon}(\mathbf{u}_{\epsilon})}{\prod_{i=1}^I q_i(u_i)^{(d_i-1)}}, \right\}, \quad (1.4)$$

où d_i est le degré de la variable u_i (*i.e.*, le nombre de facteurs f_{ϵ} dépendant de la variable u_i).

Ces deux approximations sont bien connues de la littérature et présentent l'intérêt d'être calculables par des procédures efficaces (algorithme Variational Bayes Expectation-Maximization (VBEM) pour la première, algorithme de type propagation de message pour la seconde). Pour une présentation plus détaillée mais néanmoins succincte de ces approximations, je renvoie le lecteur au chapitre 3 du manuscrit de thèse de Guillaume Beaumont [2].

1.2 Contributions et plan du document

Le cadre bayésien et l'approche analytique que permettent les approximations variationnelles sont un atout pour la résolution de nombreux problèmes inverses. J'ai ainsi travaillé et pu montrer leur intérêt dans deux grands types de problèmes :

- les décompositions parcimonieuses (structurées),
- la reconstruction de phase (informée).

Chacun de ces deux grands problèmes fait l'objet d'un chapitre du présent manuscrit. J'ai choisi d'organiser ces chapitres en deux parties :

- une première partie résumant les contributions algorithmiques apportées et pointant vers l'article qui en a résulté (reproduit en fin de manuscrit),
- une seconde partie focalisant sur une ou deux applications concrètes de ces algorithmes et pointant vers un ou deux articles résultants (reproduits en fin de manuscrit),

Ces parties sont volontairement succinctes pour ne pas nuire à la linéarité de mon propos, les articles sélectionnés permettant ensuite une lecture plus approfondie de mes contributions. Un dernier chapitre conclut mon manuscrit sur l'orientation de mes recherches futures.

1.3 Notations

Dans tout le document, les notations mathématiques suivantes sont adoptées :

- les vecteurs et matrices sont indiqués par des lettres grasses, respectivement minuscules (par exemple \mathbf{u}) et majuscules (par exemple \mathbf{D}),
- les quantités scalaires sont indiquées par des lettres minuscules (par exemple u_i),
- l'unité imaginaire est notée i .

Deux types de références bibliographiques jalonnent le manuscrit :

- les références précédées d'une lettre (par exemple [R3], [A15]) correspondent à mes contributions et sont à retrouver p.51, Annexe B.
- les références numérotées de [1] à [42] sont issues de la littérature et sont à retrouver p.41, Bibliographie.

Chapitre 2

Décompositions parcimonieuses

Les décompositions parcimonieuses cherchent à décrire un signal comme la somme pondérée d'un petit nombre de signaux élémentaires. Nées sans dire leur nom dans la communauté statistique des années 70, elles ont vu leur popularité croître avec les transformées en ondelettes [24] dans les années 90 puis l'avènement de l'acquisition compressée (compressed sensing, [13]) et la promesse (fort bien vendue) du dépassement des limites édictées par le théorème de Shannon-Nyquist. Sans parler de révolution, les décompositions parcimonieuses ont de fait prouvé leur utilité dans de nombreuses applications en traitement du signal et communication numérique où les signaux sont souvent "naturellement" parcimonieux dans des domaines de représentation que l'on connaît ou que l'on s'attache à définir. Par exemple, en traitement d'image, l'*a priori* parcimonieux a montré son intérêt pour la compression, l'inpainting (estimation de données manquantes) ou le débruitage, avec des applications dans de nombreux domaines de l'ingénierie (géophysique, imagerie biomédicale, etc.). Il n'est donc pas surprenant que beaucoup de contributions s'intéressent à ce type de décompositions.

2.1 Contribution méthodologique

Le problème attaché aux représentations parcimonieuses consiste à rechercher la représentation de dimension N , notée \mathbf{z} , d'un signal \mathbf{y} de dimension $M < N$ dans un dictionnaire \mathbf{D} (formalisé par une matrice de dimension $M \times N$), sous l'hypothèse ou contrainte que cette représentation présente peu d'éléments non nuls. Formellement, la parcimonie d'un vecteur \mathbf{z} est mesurée par la pseudo-norme ℓ_0 :

$$\|\mathbf{z}\|_0 \triangleq \text{Card}\{j | z_j \neq 0\} \quad (2.1)$$

où z_j est le j -ième coefficient dans \mathbf{z} . On peut alors envisager plusieurs formulations pour ce problème de représentations parcimonieuses : recherche de la représentation la plus parcimonieuse, celle menant à une reconstruction exacte, ou celle contrainte sur un certain niveau de parcimonie menant à l'erreur de reconstruction la plus petite... L'une des plus populaires peut-être met en jeu une régularisation implicite :

$$\hat{\mathbf{z}} = \arg \min_{\mathbf{z}} \|\mathbf{y} - \mathbf{D}\mathbf{z}\|_2 + \lambda \|\mathbf{z}\|_0 \quad (2.2)$$

où λ est un facteur lagrangien, pondérant l'importance donnée à la contrainte de parcimonie au regard du terme d'attache aux données.

Ce problème – sous-déterminé – est de complexité combinatoire. Un gros enjeu de la communauté de signal dans les années 2000 a donc été de proposer des algorithmes – sous-optimaux – pour résoudre ce problème.

La pseudo-norme ℓ_0 n'étant pas convexe, elle est souvent remplacée par la norme ℓ_1 pour laquelle on montre l'identité des solutions du problème posé avec la pseudo-norme ℓ_0 sous certaines conditions [16]. La convexité de la norme ℓ_1 permet le recours à des méthodes d'optimisation quadratique ou linéaire (selon la formulation du problème). De fait, on distingue dans la littérature deux grandes familles d'algorithmes déterministes de décomposition parcimonieuse :

- les algorithmes de poursuite, utilisant la pseudo-norme ℓ_0 , et approchant la solution parcimonieuse par une succession de décisions locales sur le support de la décomposition (*i.e.* les atomes utilisés dans la décomposition) (*cf. e.g.*, [25, 27]),
- les algorithmes basés sur la norme ℓ_1 et s'attachant à la résolution d'un problème relaxé (*cf. e.g.*, [7, 36]).

Tels quels, ces algorithmes ne peuvent prendre en compte aucun *a priori* supplémentaire sur le dictionnaire \mathbf{D} . Ainsi, ils considèrent implicitement que les atomes du dictionnaire ont tous la même probabilité d'occurrence (*i.e.* d'être choisis dans la décomposition parcimonieuse) et ne peuvent modifier cet *a priori*, se privant d'une information utile à la bonne reconstruction de la représentation parcimonieuse. Ils sont par ailleurs très mal adaptés à l'apprentissage des paramètres du problème. En particulier, tous les signaux n'ont pas une représentation parcimonieuse exacte dans un dictionnaire donné, on en calcule une approximation. Celle-ci suppose l'ajout d'un bruit sur la représentation parcimonieuse, bruit qui peut conditionner les performances de l'algorithme (en termes de qualité de reconstruction et de convergence). Ces défauts ont motivé l'utilisation de modèles bayésiens, permettant d'une part d'introduire naturellement des informations *a priori* (par exemple sur les probabilités d'occurrence des atomes) et d'autre part d'explicitier et d'estimer facilement les paramètres du problème.

Si le modèle gaussien est généralement considéré pour les observations – soit $p(\mathbf{y}|\mathbf{z}) = \mathcal{N}(\mathbf{D}\mathbf{z}, \sigma_n^2)$ – différents modèles ont été proposés dans la littérature pour formaliser l'*a priori* parcimonieux sur \mathbf{z} . L'un d'entre eux, le modèle Bernoulli-gaussien s'est avéré particulièrement pertinent. On en trouve là encore plusieurs variantes. Celle à laquelle je me suis en particulier intéressée définit les éléments du vecteur parcimonieux \mathbf{z} comme la multiplication d'une variable gaussienne \mathbf{x} et d'une variable de Bernoulli \mathbf{s} :

$$\mathbf{z} = \mathbf{x} \odot \mathbf{s} \quad (2.3)$$

avec

$$p(\mathbf{x}) = \prod_i p(x_i) \quad \text{et} \quad p(\mathbf{s}) = \prod_i p(s_i) \quad (2.4)$$

$$p(x_i) = \mathcal{N}(0, \sigma_x^2) \quad \text{et} \quad p(s_i) = \text{Ber}(p_i) \quad (2.5)$$

où s_i (respectivement x_i) est le i -ème coefficient de \mathbf{s} (respectivement \mathbf{x}), \odot représente la multiplication terme-à-terme, $\mathcal{N}(0, \sigma_x^2)$ la loi gaussienne centrée de variance σ_x^2 et $\text{Ber}(p_i)$ la loi de Bernoulli de paramètre p_i .

Dans ce modèle Bernoulli-gaussien, les variables de Bernoulli s_i jouent le rôle d'interrupteurs dans la représentation parcimonieuse (1 lorsque la colonne de \mathbf{D} indexée par s_i contribue à la représentation de \mathbf{y} , 0 sinon). Les paramètres de Bernoulli p_i peuvent alors être fixés de façon à prendre en compte l'*a priori* disponible sur la probabilité d'occurrence d'une colonne de \mathbf{D} dans la construction des observations. On appelle classiquement le vecteur \mathbf{s} support de la représentation parcimonieuse.

2.1.1 Algorithmes gloutons revisités (travaux de thèse)

Tandis que les modèles proposés dans la littérature n'avaient jamais été, à notre connaissance, explicitement reliés à l'une des formulations originales du problème de décomposition parcimonieuse,

nous avons prouvé dans [T1] et [A5] que l'ensemble des solutions du problème d'estimation jointe au sens du Maximum A Posteriori (MAP)

$$(\hat{\mathbf{x}}, \hat{\mathbf{s}}) = \arg \max_{\mathbf{x}, \mathbf{s}} \log p(\mathbf{y}, \mathbf{x}, \mathbf{s}), \quad (2.6)$$

exploitant le modèle (2.3) - (2.5), est identique sous de faibles conditions à celui du problème (2.2). Pour résoudre ce problème d'estimation MAP, nous avons introduit quatre procédures itératives généralisant quatre algorithmes de poursuite de la littérature. Elles ont recours, comme leurs homologues déterministes, à une succession de décisions locales sur le support de la décomposition, mais, de par leur cadre bayésien, elles comblent leurs limitations : i) les probabilités d'occurrence des atomes sont prises en compte dans l'algorithme par l'intermédiaire des paramètres de Bernoulli, ii) le cadre bayésien facilite l'estimation des paramètres du modèle, ici par exemple, on implémente l'estimation de la variance du bruit gaussien modélisant l'écart entre l'observation et sa reconstruction. En outre, tandis que les algorithmes de poursuite sont des procédures "forward", *i.e.* ajoutent à chaque itération un ou plusieurs atomes à la décomposition, les algorithmes que nous avons alors proposés permettent naturellement la désélection d'atomes, s'étendant ainsi à des procédures "forward-backward".

2.1.2 Algorithmes variationnels bayésiens

Dans la continuité de ces travaux de thèse, j'ai envisagé [A7] d'exploiter ce modèle Bernoulli-gaussien dans des approches reposant sur la marginalisation de certaines variables. En particulier, l'estimation du support de la représentation parcimonieuse est charnière dans le problème parcimonieux : dès lors que les colonnes de \mathbf{D} contribuant à \mathbf{y} sont identifiées – et donc la réduction de dimension réalisée –, estimer les amplitudes \mathbf{x} correspondantes devient accessible par des procédures classiques (pseudo-inverse, descente de gradient...). Dans le cadre bayésien que nous nous sommes donné, cela peut s'exprimer comme le problème d'estimation MAP marginalisée suivant :

$$\forall i \quad \hat{s}_i = \arg \max_{s_i} \log p(s_i | \mathbf{y}) \quad (2.7)$$

avec $p(s_i | \mathbf{y}) = \sum_{\mathbf{s}_{\neq i}} \int_{\mathbf{x}} p(\mathbf{s}, \mathbf{x} | \mathbf{y}) d\mathbf{x}$, où $\mathbf{s}_{\neq i}$ représente le vecteur \mathbf{s} privé de son i -ème coefficient.

Ce problème ne peut être résolu de façon analytique. Pour le résoudre, nous avons opté pour une approximation variationnelle bayésienne : l'approximation de champ moyen et l'algorithme Variational Bayes Expectation-Maximization (VBEM) qui lui est attaché. Cette approche renforce l'intérêt du cadre bayésien : par nature, elle s'appuie sur une estimation itérative des distributions *a posteriori*, plutôt que des réalisations mêmes des variables. Ainsi, elle "affine" les distributions *a posteriori* et ne prend de décision "dure" qu'à l'issue de la procédure. Cette propriété explique les bons résultats (en terme de reconstruction du support et d'estimation des coefficients de la décomposition) de l'algorithme comparé aux autres algorithmes de la littérature, et ceci, avec une complexité raisonnable. Pour souligner cette qualité, nous avons baptisé notre approche SoBaP pour Soft Bayesian Pursuit.

2.1.3 Machine de Boltzmann

Durant ma thèse, j'ai montré dans le cadre particulier de la compression d'image que si la redondance du dictionnaire peut être intéressante du point de vue de l'approximation du signal à compresser, le coût de codage des positions des atomes correspondants peut être rédhibitoire. Nous nous sommes tournés vers une structuration du dictionnaire, sous la forme d'un ensemble de bases. Par "structuration", nous entendons le fait de restreindre le nombre de combinaisons d'atomes possibles pour décrire le vecteur recherché. Ici, la "structure" est implicite, le signal est supposé avoir une représentation

parcimonieuse dans une base parmi un ensemble donné. Dans ce contexte, je me suis intéressée à deux problèmes en particulier : l'apprentissage d'un ensemble de bases [A4] et le problème d'interpolation spatiale d'image exploitant un ensemble de bases [A6].

Mon post-doctorat à l'Institut Langevin m'a permis d'approfondir l'idée de décomposition parcimonieuse structurée : en coopération avec Laurent Daudet et l'équipe Perception de l'ENS, je me suis intéressée [R2] à la possibilité de représenter de façon très simplifiée des signaux audio tout en garantissant leur identifiabilité immédiate par des auditeurs (l'idée d'“esquisse auditive”). Dans ce cadre, j'ai introduit l'usage des “machines de Boltzmann” qui constituent un modèle probabiliste plus général permettant de prendre en compte un très grand nombre de structures entre atomes.

Formellement, le modèle multiplicatif (2.3) est conservé, et les amplitudes \mathbf{x} restent supposées gaussiennes. En revanche, le modèle de Bernoulli est remplacé par une machine de Boltzmann selon :

$$p(\mathbf{s}) \propto \exp(\mathbf{b}^T \mathbf{s} + \mathbf{s}^T \mathbf{W} \mathbf{s}) \quad (2.8)$$

où \propto représente la proportionnalité, \mathbf{b} est un paramètre de biais jouant le rôle des paramètres de Bernoulli dans le modèle précédent et \mathbf{W} caractérise les interactions entre les éléments de \mathbf{s} . Cette dernière matrice permet la description de nombreux types de liens, entre des atomes par exemple très distants. On souligne ainsi [R1] qu'une machine de Boltzmann couvre un grand nombre d'autres modèles probabilistes comme le modèle non-structuré Bernoulli-gaussien, ou le modèle par chaîne de Markov. Exploitant ce modèle, j'ai proposé deux algorithmes de décompositions parcimonieuses [C1, A8, R1]. L'un d'eux [C1, R1] utilise une approximation de champ moyen, et généralise ainsi l'algorithme SoBaP introduit dans [A7]. L'autre [A8] peut être considéré comme une extension probabiliste et structurée de l'algorithme OMP (Orthogonal Matching Pursuit) [27]. Les deux algorithmes présentent l'intérêt commun de s'adapter à une très grande variété de modèles de structures parcimonieuses (de par la machine de Boltzmann) et offrent des qualités propres : décisions “douces” pour le premier (avec une estimation itérative de distributions) permettant d'obtenir de très bonnes performances en terme notamment de reconstruction du support, et très faible complexité pour l'autre. Nous avons baptisé l'extension de SoBaP à la machine de Boltzmann SSoBaP pour Structured Soft Bayesian Pursuit.

2.1.4 Article sélectionné

Les deux algorithmes SoBaP et SSoBaP ont été valorisés dans un article de la revue IEEE *transactions on signal processing* en 2012 [R1]. Cet article est reproduit en p.58 du présent manuscrit. L'état de l'art qui y est exposé reflète assez bien, je pense, celui de sa parution : SoBaP est ainsi comparé à l'algorithme EMBGAMP (pour EM Bernoulli-gaussian Approximate Message Passing), reposant sur une approximation de Bethe et une procédure par propagation de message approché (AMP), proposé dans [38].

Les procédures par propagation de message approché ont, depuis, injustement – mais suis-je vraiment objective ? – éclipsé l'approximation de champ moyen sur laquelle reposent SoBaP et SSoBaP. Il faut rappeler ici les principales différences entre ces deux soeurs d'une même famille :

- L'approximation de champ moyen contraint l'approximation de la distribution *a posteriori* cible à l'indépendance de ses variables (ou groupes de variables). La procédure VBEM qui met en œuvre la résolution du problème de minimisation (1.1) est itérative, mettant à jour l'ensemble des facteurs de l'approximation *successivement* à chaque itération.
- L'approximation de Bethe propose une approximation moins simplificatrice mais plus complexe et de fait plus coûteuse en calculs que l'approximation de champ moyen. L'algorithme par propagation de message approché qui a été proposé pour résoudre la minimisation (1.1) répond

à ce surcoût par une simplification des facteurs de l'approximation et une mise à jour *parallèle* de ces facteurs. Cette simplification supplémentaire est indolore dans le cas de l'acquisition compressée pour lequel AMP a été proposé : AMP présente une bonne convergence pour des matrices \mathbf{D} dont les éléments sont indépendants et identiquement distribués de moyenne nulle [1, 9, 21]. De fait, les diagrammes de phase proposés pour illustrer les performances des différentes approches de la littérature montrent bien la supériorité des approches AMP dans ces scénarios (*cf.* Fig. 2.1). En revanche, dès lors qu'on s'éloigne de ces scénarios simples, ses performances chutent drastiquement [3].




Figure 2.1 – Diagramme de phase (MSE normalisée en fonction de $\alpha = M/N$ et $\rho = K/N$ avec K nombre de coefficients non nuls dans \mathbf{z}) pour les performances de SoBaP avec un bruit additif de variance $\sigma_n^2 = 10^{-8}$, $N = 1024$, et \mathbf{D} matrice d'éléments gaussiens *i.i.d.* de moyenne nulle et de variance unité. Les lignes indiquent (de bas en haut) le seuil optimal pour l'acquisition compressée comprimée [40] (ligne droite en pointillés), la transition de phase AMP tel qu'implémenté dans [21] (ligne continue grasse) et la transition de Donoho-Tanner pour la reconstruction convexe ℓ_1 [7] (ligne continue fine). (extrait de [20])

Ainsi, si les algorithmes AMP sont souvent retenus comme références des algorithmes bayésiens, il ne faut pas oublier que leur utilisation dans le cadre des décompositions parcimonieuses est fortement assujettie à la nature du dictionnaire \mathbf{D} , là où les approximations de champ moyen se montreront plus robustes.

2.2 Application : estimation modale en acoustique sous-marine (ultra) basse fréquence

A mon arrivée à l'ENSTA Bretagne, j'ai focalisé mon attention sur le signal acoustique sous-marin. En particulier, deux grands problèmes m'ont occupée (et m'occupent encore) : l'estimation modale en ultra-basses fréquences et milieu petit-fond, et l'estimation de directions d'arrivée en milieu fluctuant. Le premier a été l'occasion pour moi de confronter SoBaP à la "vraie" vie des données.

En acoustique sous-marine, les environnements petit-fond (d'une profondeur de quelques centaines de mètres) agissent comme des guides d'onde dispersifs lorsqu'on considère des sources de basses fréquences (inférieures à quelques centaines de hertz). Les signaux propagés sont alors classiquement décrits comme la somme de composantes dites "modales" se propageant longitudinalement selon leurs

nombre d'onde horizontaux. On peut ainsi décrire le signal reçu à la fréquence f et à la distance r de la source comme [19]

$$y(f, r) = Q \frac{s(f)}{\sqrt{r}} \sum_{m=1}^{M(f)} \psi_m(f, z_s) \psi_m(f, z) \frac{e^{-i r k_{rm}(f)}}{\sqrt{k_{rm}(f)}} + n(f, r) \quad (2.9)$$

où $i^2 = -1$, Q est une constante multiplicative, $n(f, r)$ un bruit additif, $M(f)$ le nombre de modes se propageant à la fréquence f , $k_{rm}(f)$ le nombre d'onde attaché au m -ième mode propagatif à la fréquence f et $\psi_m(f, z)$ sa fonction modale en profondeur. La connaissance de ces composantes modales est d'une grande importance pour la caractérisation du milieu d'observation et, en conséquence, pour la localisation de sources [19].

2.2.1 SoBaP informé

Selon le contexte opérationnel et l'antenne à disposition, plusieurs outils sont envisageables pour résoudre le problème d'estimation des composantes modales. Dans le cas d'une antenne linéaire horizontale, l'estimation modale équivaut à une estimation spectrale (les nombres d'ondes horizontaux jouant le rôle de pulsations spatiales, il s'agit finalement de considérer une transformée de Fourier spatiale sur la longueur de l'antenne). Plus formellement, le modèle physique (2.9) peut se réécrire de façon matricielle et discrétisée comme

$$\mathbf{y}_\nu = \mathbf{D} \mathbf{z}_\nu + \mathbf{n}_\nu \quad (2.10)$$

où $\mathbf{y}_\nu \in \mathbb{C}^M$ contient les pressions acoustiques reçues sur l'antenne à l'indice fréquentiel ν (tel que $f = \nu \Delta_f$ avec Δ_f pas d'échantillonnage fréquentiel), $\mathbf{D} \in \mathbb{C}^{M \times N}$ est une matrice de Fourier (spatiale), et $\mathbf{z}_\nu \in \mathbb{C}^N$ contient les amplitudes modales des différents modes se propageant à l'indice fréquentiel ν . L'estimation modale revient donc à estimer \mathbf{z}_ν selon les observations \mathbf{y}_ν . On note qu'en l'absence de la connaissance du nombre de modes propagatifs attendus, une recherche sur une grille telle que $M > N$ peut être considérée. Il s'agit alors d'ajouter des hypothèses sur l'inconnue \mathbf{z}_ν de façon à réduire l'espace de recherche.

On peut dès lors envisager nombre d'outils bien connus de la littérature. Ces méthodes sont cependant souvent assujetties à des contraintes sur l'acquisition (échantillonnage régulier des capteurs) ou supposent des hypothèses fortes sur la nature du milieu de propagation ou des sources envisagées (stationnarité). Afin d'alléger ces limitations, il est raisonnable de s'appuyer sur la connaissance physique de la propagation. Deux informations sont en particulier d'intérêt :

- peu de modes se propagent, appelant ainsi naturellement l'hypothèse de parcimonie sur \mathbf{z}_ν ,
- en présence de sources large bande, il existe une expression explicitant le comportement des modes à travers les fréquences en jeu : on peut en effet montrer que

$$k_{rm}[\nu + 1]^2 = k_{rm}[\nu]^2 + (2\nu + 1) \left(\frac{2\pi\Delta_f}{c} \right)^2 + \epsilon[\nu] \quad (2.11)$$

où c est la vitesse de l'onde acoustique dans le milieu et $\epsilon[\nu]$ explicite la dépendance (faible) aux nombres d'ondes verticaux.

Dans [R8], j'ai proposé d'adapter SoBaP à ce problème particulier. Ainsi, l'approche intègre une particularisation des paramètres du modèle bayésien dans une procédure itérative permettant de suivre les modes propagatifs à travers la bande de fréquences de la source. La prise en compte de la relation de dispersion permet de réduire le nombre d'observations – de capteurs – nécessaires à l'estimation modale sur l'ensemble des fréquences de la source. La figure 2.2 illustre, par des diagrammes

fréquences-nombres d'ondes obtenus sur données réelles, les résultats de l'approche proposée au regard de l'approche classique consistant à réaliser une transformée de Fourier spatiale inverse et une approche considérant uniquement un a priori parcimonieux, comme OMP [27].




Figure 2.2 – Reconstruction des diagrammes fréquences-nombres d'onde à partir de données acquises par 30 capteurs avec un rapport signal à bruit de 10 dB : (a) par simple transformée de Fourier inverse, (b) par l'algorithme OMP, (c) par l'algorithme SoBaP proposé. (extrait de [N2])

2.2.2 Travaux ultérieurs et concomitants

Thèse de Thomas Paviet-Salomon

Dans la continuité de ces travaux, nous avons pu obtenir un financement de thèse de l'AID et de l'ONRG (Office of Naval Research Global). Cette thèse – menée par Thomas Paviet-Salomon – a été l'occasion d'aborder la problématique sous l'angle de la parcimonie dans un dictionnaire continu. Plus précisément, elle a permis d'étudier une approche intégrant une descente de gradient dans un algorithme de poursuite (OMP). Ses performances ont pu être à la fois quantifiées sur données synthétiques et jugées qualitativement sur deux jeux de données réelles, résultant en une analyse empirique assez complète de l'approche [A26, R12].

Post-doctorat de Clément Dorffer

Dans une continuité plus méthodologique de mes travaux, le post-doctorat de Clément Dorffer a également permis de débiter des premières recherches sur la combinaison entre l'apprentissage des modèles physiques sous-tendant la propagation acoustique en milieu petit fond et leur prise en compte dans la résolution de l'estimation modale dans un cadre bayésien. Pour ce faire, nous nous sommes intéressés [A27] à un autre modèle de machine de Boltzmann, la machine de Boltzmann restreinte, telle que :

$$p(\mathbf{s}) = \sum_{\mathbf{h}} p(\mathbf{s}, \mathbf{h}) \propto \exp(\mathbf{b}^T \mathbf{s} + \mathbf{s}^T \mathbf{W} \mathbf{h} + \mathbf{a}^T \mathbf{h}) \quad (2.12)$$

où \propto représente la proportionnalité, \mathbf{b} un paramètre de biais attaché au support \mathbf{s} de la représentation parcimonieuse, et \mathbf{h} joue le rôle d'intermédiaire structurant les liens entre les éléments de \mathbf{s} par la matrice \mathbf{W} . Le machine de Boltzmann restreinte (RBM) est au cœur de nombreuses contributions en apprentissage automatique mettant l'accent sur deux qualités précieuses : i) elle a été identifiée comme

modèle générique pouvant approcher n'importe quelle distribution sur $\{0, 1\}^N$ [32]; ii) des algorithmes efficaces ont été développés pour l'entraîner, ce qui en fait un puissant modèle de représentation lorsque de grandes quantités de données sont disponibles (e.g., [18]). Apprise sur des bases de données simulées (selon des guides de Pekeris [19] dans un premier temps), la RBM permet de fait une bonne représentation de la relation (2.11) liant les nombres d'ondes des modes propagatifs d'une fréquence à l'autre. Intégrée dans un second temps dans un algorithme variationnel bayésien exploitant une approximation de champ moyen, elle résulte ainsi en une estimation plus précise des nombres d'onde recherchés.

Collaboration avec Cédric Herzet

Les dictionnaires continus que l'on trouve notamment dans ces deux dernières contributions ont fait l'objet d'une collaboration particulière avec Cédric Herzet (INRIA Centre Rennes). Les algorithmes de décomposition parcimonieuse sur dictionnaires continus peuvent être grossièrement divisés en deux catégories :

- les algorithmes comme celui proposé dans [41] sont basés sur la formulation et la résolution d'un problème d'optimisation convexe particulier, appelé "semi-definite programming" et focalisent généralement sur un dictionnaire particulier (par exemple un dictionnaire de Fourier),
- les algorithmes dérivés de l'approche par descente de gradient reposent notamment sur une étape dite de "sélection d'atome", dans laquelle on sélectionne la colonne d'un dictionnaire grossier la plus corrélée aux observations (dans notre cas à la pression acoustique reçue le long de l'antenne). Une étape de descente de gradient est alors effectuée afin de raffiner la sélection.

La première famille d'algorithmes est en soi très coûteuse du point de vue de la complexité algorithmique et peut difficilement être améliorée dans la mesure où la méthode d'optimisation sur laquelle elle repose fait appel à des calculs matriciels de grande dimension. La seconde famille d'algorithmes est peut-être la plus intuitive. C'est celle que nous avons exploitée dans la thèse de Thomas Paviet-Salomon et dans le post-doctorat de Clément Dorffer. Outre sa simplicité de mise en oeuvre, on peut espérer en réduire la complexité algorithmique par les méthodes de "screening". Parmi ces dernières, nous nous sommes intéressés en particulier à deux d'entre elles :

- la méthode de "dual screening" propose une procédure permettant de diminuer le coût de calculs de certaines opérations, notamment la "sélection d'atome" mentionnée ci-dessus [A25],
- la méthode de "joint screening" s'attache à transformer le problème initial de grande dimension en un problème ayant la même solution, mais de dimension inférieure [R11, A21].

2.2.3 Articles sélectionnés

L'intégration de la relation de dispersion (2.11) dans la procédure SoBaP a donné lieu à un article de la revue *IEEE journal of oceanic engineering* en 2015 [R8]. Il est reproduit en p.72 de ce manuscrit.

L'algorithme DSSoBaP élaboré avec mon post-doctorant Clément Dorffer a fait l'objet d'un article dans les actes de la *European Signal Processing Conference (EUSIPCO)* en 2021 [A27]. Il est reproduit en p.81 de ce manuscrit. Nous travaillons actuellement à la rédaction d'un article de journal.

2.3 Synthèse et ouverture

L'estimation parcimonieuse a constitué ma voie d'entrée dans les algorithmes variationnels bayésiens.

En questionnant d'abord les modèles propices à la bonne description de l'hypothèse parcimonieuse puis des structures attendues dans les décompositions, j'ai pu prendre la mesure de ce qui font à mes yeux la beauté et la force du cadre bayésien : l'explicitation des hypothèses posées sur les inconnues.

Le problème d'estimation posé, il est toujours possible ensuite de revenir au cadre déterministe. Mais ce serait alors se priver de la prise en compte de l'incertitude – du doute – permise par les outils probabilistes. Là encore, cette incertitude est cadrée par le modèle mais elle est quantifiée et contrôlée.

Pour m'y intéresser depuis peu, le cadre bayésien permet également une compréhension plus aisée de l'apprentissage automatique, ou du moins lui donne-t'il un éclairage qui m'aide à y trouver du sens. La question de l'estimation des hyperparamètres, par exemple, au coeur de la problématique bayésienne, sensibilise de fait à la dichotomie "modèle v.s. données" : où s'arrête le modèle ? Où commence l'apprentissage ? Les machines de Boltzmann (restreintes) que j'ai introduites en parcimonie structurée permettent d'élaborer des algorithmes génériques, capables de considérer n'importe quelle structure. Mais cette structure est dans les paramètres de la machine et dans l'apprentissage. Vaut-il mieux alors un modèle mal adapté dans sa construction mais bien contrôlé (avec moins d'hyperparamètres par exemple), ou un modèle générique mais mal appris (parce que peu de données par exemple) ?

Chapitre 3

Reconstruction de phase

La reconstruction de phase, ou plus exactement le problème visant à reconstruire un vecteur complexe en ne tenant compte que de l'amplitude des mesures, intéresse de nombreux domaines applicatifs. Durant mon post-doctorat au laboratoire de physique statistique de l'ENS, j'ai eu la chance d'y être confrontée sur une belle application d'imagerie optique en collaboration avec le laboratoire Kastler Brossel. Arrivée à l'ENSTA Bretagne, j'ai pu m'appuyer sur mes contributions pour proposer une extension de ces travaux à la reconstruction de phase "informée", formalisation possible de l'estimation de directions d'arrivée en milieu océanique fluctuant.

3.1 Contribution méthodologique

Dans le problème de reconstruction de phase "classique", les observations sont privées de toute information de phase. Formellement, le problème s'écrit comme suit : observant $\mathbf{y} \in \mathbb{R}_+^M$ par une matrice de mesures complexe \mathbf{D} de dimensions $M \times N$, on cherche à retrouver le signal $\mathbf{z} \in \mathbb{C}^N$ tel que

$$\mathbf{y} = |\mathbf{Dz} + \mathbf{n}| \quad (3.1)$$

Ce problème est non-convexe, notoirement difficile à résoudre. Comme en décompositions parcimonieuses, la littérature s'est alors attachée à proposer de nombreux algorithmes pour approcher des solutions de ce problème. On peut grossièrement les diviser en deux grandes familles :

- Les algorithmes par projections alternées [14, 11, 17] font succéder des projections dans l'espace image de la matrice de mesures et dans l'espace où vit le signal recherché.
- Les algorithmes basés sur les relaxations convexes [4, 39] proposent d'approcher le problème (3.1) par des problèmes relaxés qui peuvent être résolus efficacement par des procédures d'optimisation standard.

Outre ces approches, des contributions plus récentes se sont intéressées à la reconstruction "compressée" de phase. L'idée est de s'appuyer sur une hypothèse supplémentaire – le signal recherché est parcimonieux dans une base connue – pour diminuer le nombre de mesures nécessaires à la résolution du problème. Les travaux s'intéressant à cette approche sont encore peu nombreux, on retrouve les grandes familles citées ci-dessus. Les procédures d'optimisation convexes sont à la base des contributions [23, 26], tandis que des approches par optimisations locales successives sont par exemple proposées dans [33].

Le point de vue bayésien pour la reconstruction de phase a été abordé par le biais de l'acquisition compressée dans [34]. Modélisant l'absence de phase par une variable cachée, l'approche proposée

repose sur une approximation variationnelle de Bethe et fait appel à un algorithme de type somme-produit. Alternative à cette approche, nous avons proposé dans [A15] une autre méthode de résolution, basée sur un algorithme variationnel bayésien Expectation-Maximisation (VBEM), exploitant une approximation de champ moyen.

Le formalisme adopté est le même dans les deux cas. Chaque mesure réelle positive y_k est exprimée, $\forall k \in \{1, \dots, M\}$, comme

$$y_k = e^{i\theta_k} \left(\sum_{i=1}^N z_i d_{ki} + n_k \right) = e^{i\theta_k} ((\mathbf{Dz})_k + n_k) \quad (3.2)$$

où $i^2 = -1$, $\theta_k \in [0, 2\pi[$ représente le conjugué complexe de sa phase manquante, d_{ki} est le k -ème élément de la i -ème colonne \mathbf{d}_i du dictionnaire \mathbf{D} , $(\mathbf{Dz})_k$ est le k -ème coefficient du vecteur issu du produit matriciel de \mathbf{D} et \mathbf{z} , et n_k est un bruit circulaire gaussien de moyenne nulle et de variance σ_n^2 . En l'absence d'information plus précise sur la phase, on peut la supposer distribuée uniformément sur $[0, 2\pi[$, soit

$$p(\boldsymbol{\theta}) = \prod_{k=1}^M p(\theta_k) \quad \text{avec} \quad p(\theta_k) = \frac{1}{2\pi} \quad \forall k \in \{1, \dots, M\}. \quad (3.3)$$

On note que sous ces hypothèses, l'absence d'observation des phases est naturellement prise en compte dans le modèle puisque la marginalisation sur θ_k conduit à une distribution sur y_k qui dépend seulement des modules de y_k et $(\mathbf{Dz})_k$:

$$p(y_k|\mathbf{z}) \propto I_0 \left(\frac{2}{\sigma_n^2} |y_k^* (\mathbf{Dz})_k| \right) \exp \left(-\frac{1}{\sigma_n^2} \left(|y_k|^2 + |(\mathbf{Dz})_k|^2 \right) \right) \quad (3.4)$$

où \propto représente la proportionnalité, $*$ le conjugué complexe et $I_0(\cdot)$ la fonction de Bessel modifiée de première espèce à l'ordre 0.

Selon l'information *a priori* dont nous disposons sur \mathbf{z} , plusieurs distributions peuvent être envisagées. Dans le problème de reconstruction de phase "classique", il est admis de considérer un *a priori* gaussien :

$$p(\mathbf{z}) = \prod_{i=1}^N p(z_i) \quad \text{avec} \quad p(z_i) = \mathcal{CN}(0, \sigma_x^2) \quad (3.5)$$

C'est le modèle adopté pour le développement algorithmique que j'ai proposé dans [A15]. Cet *a priori* peut être remplacé par un modèle Bernoulli-gaussien (2.3)-(2.5) pour prendre en compte une hypothèse de parcimonie sur \mathbf{z} , comme j'ai pu le considérer dans [R3].

Selon le modèle considéré, le problème de reconstruction de phase peut être exprimé de manière différente. Dans [A15], nous proposons de nous attacher à la résolution du problème d'estimation MAP marginalisé :

$$\hat{\mathbf{z}} = \arg \max_{\mathbf{z}} \log p(\mathbf{z}|\mathbf{y}) \quad (3.6)$$

avec $p(\mathbf{z}|\mathbf{y}) = \int_{\boldsymbol{\theta}} p(\mathbf{z}, \boldsymbol{\theta}|\mathbf{y}) d\boldsymbol{\theta}$. En raison de la marginalisation sur les variables cachées $\boldsymbol{\theta}$, le calcul direct de $p(\mathbf{z}|\mathbf{y})$ est irréaliste. Dans [34], la distribution $p(\mathbf{z}, \boldsymbol{\theta}|\mathbf{y})$ est approchée par une approximation de Bethe et un algorithme AMP est considéré. Dans [A15], nous leur avons préféré une approximation de champ moyen et une procédure VBEM. Nous avons ainsi montré sur des résultats préliminaires et simples que notre approche offrait des qualités désirables, notamment en terme de temps de calculs et de robustesse au bruit (cf. Fig. 3.1).




Figure 3.1 – Corrélation moyennée sur 100 tests, entre les signaux estimés $\hat{\mathbf{z}}$ et les signaux originaux \mathbf{z}_0 ayant généré les observations \mathbf{y} (avec \mathbf{D} matrice d'éléments gaussiens *i.i.d.* de moyenne nulle et de variance unité), en fonction du nombre de mesures M (l'axe des abscisses est M/N avec $N = 64$) : cas sans bruit (a), $\sigma_n^2 = 0,3$ (b) et $\sigma_n^2 = 0,7$ (c). Comparaison entre l'algorithme proposé dans [A15] appelé prVBEM, et les approches prGAMP, PhaseCut et Gerchberg-Saxton, respectivement proposées dans [34], [39] et [14]. (extrait de [A15])

3.2 Reconstruction aveugle : propagation d'ondes optiques en milieux complexes

Pendant mon post-doctorat à l'ENS, j'ai eu la chance d'aborder le problème de reconstruction de phase par le biais de deux cas pratiques de la propagation d'ondes optiques à travers des milieux complexes : la calibration du milieu et le contrôle du front d'onde permettant la focalisation.

La propagation d'ondes en milieux complexes est un problème fondamental en physique. En optique, lorsque la lumière passe au travers d'un milieu fortement diffusant, tel un tissu biologique, elle est rapidement atténuée, empêchant les techniques d'imagerie conventionnelle ; de plus le caractère aléatoire des diffusions génère un bruit dit de speckle, généralement considéré comme gênant pour l'imagerie. Dans ce contexte, le contrôle des fronts d'onde par des modulateurs spatiaux de lumière, et plus particulièrement des matrices de micro-miroirs (DMD), s'est imposé comme l'outil permettant de manipuler une lumière cohérente diffusée. Parmi les méthodes exploitant ces DMD, une approche consiste à mesurer dans un premier temps la matrice dite de transmission, caractérisant la propagation de la lumière à travers le milieu, du modulateur au détecteur. Ceci pose cependant plusieurs difficultés pratiques. D'une part, les capteurs ne mesurent que l'intensité du champ complexe. Par ailleurs, les DMD n'autorisent que des modulations d'amplitude binaires (*i.e.* ON ou OFF sur chaque miroir). Enfin, le bruit expérimental relativement important rend difficile l'estimation. La figure 3.2 schématise l'expérience menée. Le dispositif expérimental est simple : une source de lumière cohérente (laser), modulée spatialement en amplitude par le DMD, passe à travers un milieu diffusant (ici une simple couche de peinture blanche), avant d'être captée sur une caméra CCD. Les détails expérimentaux sont décrits dans l'article [R3] reproduit en p.86.

3.2.1 Calibration du milieu diffusant

Le modèle sous-jacent de la propagation d'ondes en milieux complexes est linéaire : le signal \mathbf{y} en sortie, pour une image \mathbf{z} affichée sur le DMD, est alors tel que $\mathbf{y} = |\mathbf{D}\mathbf{z}|$ avec $\mathbf{y} \in \mathbb{R}_+^M$ et où \mathbf{D} , matrice de transmission de dimension $M \times N$, est à valeurs complexes, supposées réalisations de




Figure 3.2 – Schéma expérimental : un faisceau laser est dilaté à travers un télescope afin d'obtenir un faisceau homogène. A travers un masque rectangulaire, il illumine le DMD qui agit comme modulateur spatial de lumière à amplitude binaire. Le DMD réfléchit la lumière dans deux directions différentes correspondant soit à "ON" (transmission de l'unité) soit à "OFF" (la lumière est déviée vers une décharge de faisceau). Le motif transmis est focalisé par une première lentille L1 sur le support de diffusion – ici une couche de peinture blanche –, agissant comme un support épais à diffusion multiple. Le motif de speckle transmis est collecté par un objectif de microscope et observé à travers un polariseur P sur une caméra CCD. (extrait de [R3])

variables aléatoires gaussiennes *i.i.d.* lorsque la diffusion est suffisamment forte.

L'estimation de la matrice de transmission s'exprime comme un problème de calibration supervisée : on cherche \mathbf{D} connaissant un ensemble de P couples $\{(\mathbf{z}_p, \mathbf{y}_p)\}_{p \in \{1, \dots, P\}}$. Formellement, le problème se réduit, par transposition, à une estimation ligne par ligne de la matrice \mathbf{D} , en absence de phase, selon la résolution du problème inverse :

$$\mathbf{Y}^T = |\mathbf{Z}^H \mathbf{D}^H| \quad (3.7)$$

où $\mathbf{Y} \triangleq [\mathbf{y}_1, \dots, \mathbf{y}_P] \in \mathbb{R}_+^{M \times P}$, $\mathbf{Z} \triangleq [\mathbf{z}_1, \dots, \mathbf{z}_P] \in \mathbb{C}^{N \times P}$, \cdot^H représente la transposition de la matrice complexe conjuguée et $|\cdot|$ est le module appliqué ici élément-par-élément à la matrice qu'il encadre. Ainsi formulé, le modèle révèle un problème de reconstruction de phase "classique" : étant donnée la matrice des entrées \mathbf{Z}^H , chaque colonne de \mathbf{Y}^T est utilisée pour estimer chaque colonne à valeur complexe de \mathbf{D}^H .

3.2.2 Focalisation à travers le milieu diffusant

La connaissance de la matrice de transmission peut ensuite être exploitée pour estimer la configuration du DMD permettant une observation donnée. Je me suis en particulier intéressée au cas de la focalisation de la lumière avec un maximum d'intensité sur un ou plusieurs points cibles. La figure




Figure 3.3 – Illustration de la focalisation de la lumière sur 3 points. Les cercles marquent les positions des cibles. (extrait de [N1])

3.3 illustre une sortie typiquement attendue sur la caméra pour une focalisation en trois points. Le problème se pose “classiquement” selon les termes du modèle d’observation (3.1).

La spécificité du problème vient de la nature de \mathbf{z} : les DMD n’autorisant que deux positions ON ou OFF, on les modélise ici comme des réalisations binaires de variables de Bernoulli.




Figure 3.4 – Gain en intensité expérimental ($\eta = I_{\text{foc}}/I_{\text{back}}$ avec I_{foc} l’intensité mesurée sur les cibles après optimisation de la modulation binaire, I_{back} l’intensité moyenne de l’arrière-plan) en fonction du nombre de cibles. Courbe en pointillés : meilleur fit en $1/T$. (extrait de [N1])

La figure 3.4 présente les résultats en fonction du nombre de cibles, sous la forme de boîtes. Pour chacune d’entre elles, le segment du milieu représente le gain moyen $\bar{\eta}$ sur 100 tests, les bornes supérieure et inférieure du rectangle définissent l’intervalle $[\bar{\eta} - \sigma_{\eta}, \bar{\eta} + \sigma_{\eta}]$ (où σ_{η} est l’écart-type empirique), dans lequel se trouve, sous hypothèse gaussienne, la majorité des tests, les moustaches représentent les valeurs minimum et maximum du gain, relevées sur l’ensemble des tests. La figure tend à montrer expérimentalement une dépendance en $1/T$ avec le nombre de cibles T (illustrée par la courbe en pointillés rouges), ce qui va dans le sens d’observations similaires dans la littérature [37].

3.2.3 Article sélectionné

J'ai publié le développement algorithmique proposant la particularisant de l'approche VBEM au modèle de reconstruction de phase (3.2)-(3.5) dans les actes de la conférence *IEEE international conference on acoustics, speech and signal processing (ICASSP)* de 2015 [A15]. Son application à l'optique des milieux complexes a fait l'objet d'un article de la revue *Optics Express* de la même année [R3]. C'est cet article que j'ai choisi de reproduire, en p.86 du manuscrit.

J'ai beaucoup d'attachement pour cette dernière contribution. Elle m'a donné l'opportunité de travailler avec des collègues de domaines très différents et m'a permis d'aller de l'acquisition des données jusqu'à leur traitement par un algorithme que je proposais. J'en ai été particulièrement enthousiasmée.

3.3 Reconstruction informée : estimation de DOA en présence de bruit de phase

En acoustique sous-marine passive (*i.e.* sans émission, basée uniquement sur l'écoute du milieu), les détection et localisation de sources sont fortement impactées par les fluctuations du milieu de propagation : le champ acoustique mesuré sur l'antenne réceptrice présente des écarts à celui attendu, induisant des erreurs d'estimation. Si la modélisation des différents phénomènes à l'origine de ces fluctuations et la prise en compte fine de la propagation acoustique dans ces milieux sont difficiles, on peut tout de même chercher à modéliser leurs conséquences sur les champs acoustiques mesurés afin de les intégrer dans les traitements. On sait notamment que les ondes internes [8] sont responsables de variations spatio-temporelles du profil de célérité du son et entraînent une déformation du front d'onde incident. Cette déformation peut être prise en compte via des déphasages aléatoires. Etablissant un lien entre cette problématique et mes travaux en optique des milieux complexes, j'ai proposé d'exploiter le cadre bayésien posé dans [A15, R3] pour prendre en compte la connaissance physique des déphasages attendus. Plus spécifiquement, tandis que le problème de reconstruction de phase recourt naturellement à une loi uniforme sur $[0, 2\pi[$ pour modéliser l'absence de phase des observations, il s'agit ici de considérer des distributions de probabilité plus informatives. Dans [A19, C5], c'est ainsi un modèle de Markov qui est proposé :

$$p(\boldsymbol{\theta}) = \prod_{k=2}^M p(\theta_k | \theta_{k-1}) p(\theta_1) \quad (3.8)$$

avec $p(\theta_k | \theta_{k-1}) = \mathcal{N}(a\theta_{k-1}, \sigma_\theta^2)$, $\forall k \in \{2, \dots, M\}$, $a \in \mathbb{R}_+$ et $p(\theta_1) = \mathcal{N}(0, \sigma_{\theta_1}^2)$. D'un point de vue pratique, ce modèle permet de décrire les fluctuations spatiales du milieu de propagation comme déformation du front d'onde le long de l'antenne ; l'intensité de ces fluctuations est liée à la valeur du paramètre σ_θ^2 .

La procédure proposée pour exploiter ce modèle reprend les grandes lignes de l'algorithme de reconstruction de phase introduit dans [A15, R3] (moyennant quelques calculs supplémentaires et une approximation mettant en jeu une distribution de Von Mises).

3.3.1 Travaux ultérieurs et concomittants

Thèse de Guillaume Beaumont

Dans la continuité de ces travaux préliminaires, j'ai obtenu un financement de thèse de la DGA. Cette thèse, menée par Guillaume Beaumont, a été l'occasion d'aborder une autre approximation variationnelle bayésienne, celle de Bethe, et de la particulariser au problème d'intérêt [A22], ainsi que

de considérer un autre modèle probabiliste, celui de la distribution de Von Mises multidimensionnelle, distribution directionnelle naturellement plus adaptée à la modélisation d'un bruit de phase [A23]. Les algorithmes développés ne sont encore qu'au stade de la preuve de concept et mériteraient que l'on se penche sur une confrontation à des données réelles.

Post-doctorat de Riwal Lefort

En parallèle de ces travaux, j'ai encadré un post-doctorat de trois ans sur la thématique d'estimation de directions d'arrivée en milieu océanique fluctuant (projet de l'Agence Innovation Défense (AID)). Ce post-doctorat fut l'occasion d'aborder la problématique sous un angle plus pragmatique en proposant des approches de traitement par sous-antennes [A20, R9] ou encore en testant la possibilité d'appliquer des algorithmes d'apprentissage automatique [R7] sans *a priori* physique. Ces deux approches se sont conclues en demi-teinte. Le traitement par sous-antennes repose sur un traitement local qui se prive de l'information portée par la dimension de l'antenne considérée, à moins de considérer une modélisation plus complexe, permettant une propagation de l'information utile d'une sous-antenne à l'autre. De leur côté, les approches par apprentissage automatique trouvent leur intérêt lorsque de grandes bases de données sont disponibles. Ainsi, nous avons pu les envisager dans le cadre des données acquises en cuve par Gaultier Réal pendant sa thèse [31]. Cependant l'accès à de grandes bases de données est questionnable dans le contexte des milieux de propagation fluctuants. On peut l'envisager dans le cadre d'observatoires acoustiques, mais ces approches semblent bien peu réalistes dans le cas d'antennes embarquées ou remorquées. En revanche, elles ouvrent la (large) perspective des modèles hybrides physiques/appris.

3.3.2 Article sélectionné

Pour illustrer la particularisation et le développement d'un algorithme variationnel bayésien dédié à la reconstruction de phase informée, j'ai choisi de pointer vers l'article dont je suis la première autrice. Cet article est un article de conférence, publié dans les actes de la conférence *IEEE international conference on acoustics, speech and signal processing (ICASSP)* de 2017 [A19]. Je le reproduis en p.100 du manuscrit. Une version plus détaillée est disponible en rapport technique sur arXiv.

3.4 Synthèse et ouverture

Le problème de reconstruction de phase est un problème répandu qui touche grand nombre de cas pratiques. En bayésien, le recours à une variable latente pour modéliser l'absence de phase permet de contourner la non-linéarité du module. L'approximation de champ moyen à ce type de modèle se développe assez naturellement par algorithme VBEM, moyennant le passage par une distribution de Von Mises [T2].

L'estimation de directions d'arrivée en milieu océanique fluctuant interroge plus encore la non-linéarité des modèles d'observation. Le modèle de bruit de phase informé est un premier pas vers des traitements d'antenne plus intelligents, mais il reste simpliste au regard de la complexité des phénomènes océanographiques en jeu. Un vaste champ de recherche reste ici à explorer.

Chapitre 4

Perspectives de recherche

Dans ce manuscrit, j'ai fait le choix de la continuité : montrer, mettre en lumière les travaux pour lesquels j'ai "suivi mon fil", celui des algorithmes variationnels bayésiens. La recherche étant aussi affaire d'opportunités et d'affinités, j'ai eu cependant l'occasion de collaborer sur d'autres thématiques. Dans ce chapitre, je reviens brièvement sur ces travaux avant de présenter mon projet de recherche.

4.1 De ce que je n'ai pas écrit jusque là

Représentation et alignement de séquences temporelles

Mon post-doctorat à Télécom ParisTech (2011-2013) a pris place au sein du groupe AAO (maintenant ADASP) et s'est inscrit dans le double cadre européen du projet REVERIE et du réseau d'excellence 3DLIFE. Dans ce contexte coopératif, je me suis attachée à la résolution d'un nouveau problème, celui de l'analyse de mouvement par acquisition multimodale (*i.e.* s'appuyant sur un réseau de capteurs de différentes natures, ici, des caméras, Kinects, capteurs inertiels, microphones...). Mes recherches ont en particulier porté sur deux grandes problématiques : la représentation et l'alignement.

Représentation de mouvements de danse - L'idée de cet axe de recherche est de tirer parti de connaissances *a priori* sur le geste pour en proposer une décomposition en gestes élémentaires efficace. Cette approche est particulièrement pertinente en analyse de mouvements structurés comme la danse. Le point de vue adopté est bayésien : sur chacune des composantes de la décomposition, un modèle *a priori* probabiliste est posé, en processus gaussien. Cette étude prospective [A10] a de nombreuses perspectives intéressantes. A titre d'exemple, une telle décomposition peut être utile dans des tâches de reconnaissance de gestes, pour discriminer certaines caractéristiques communes à des gestes différents.

Alignement de séquences spatio-temporelles multimodales - L'estimation de correspondances entre deux séquences temporelles est une problématique partagée par de nombreuses communautés de recherche. Adoptant un point de vue à la fois multimodal et bayésien, j'ai proposé [A12] une méthode d'alignement prenant en compte deux types de données, issues, d'une part, des cartes de profondeurs acquises par des Kinects et d'autre part de microphones, et reposant sur un modèle de Markov. L'alignement obtenu par ce modèle surpasse les méthodes classiques de la littérature. En particulier, la multimodalité apporte une plus-value significative, directement vérifiable.

Parallèlement à ces travaux, j'ai eu l'occasion de collaborer avec l'équipe AAO sur des aspects purement audio, qui ont donné lieu à deux publications [A11, A13].

Modèles dynamiques imbriqués

En parallèle de mes travaux sur la reconstruction de phase en optique des milieux complexes pendant mon post-doctorat à l'ENS, j'ai eu l'opportunité de travailler avec des collègues de l'INRIA sur un autre type de modèle d'observation : les modèles dynamiques imbriqués.

Afin de prendre en compte ce type de modèle de façon efficace, nous avons proposé dans [A14] et [C2] de recourir à une procédure backward-forward bien connue de la théorie du contrôle optimal [6] pour estimer le gradient de la fonction objectif. La procédure peut ensuite être intégrée à n'importe quel algorithme de décomposition parcimonieuse basé gradient. Nous avons par exemple considéré la méthode introduite dans [5], reposant sur une technique de majoration-minoration. Particularisée au problème d'estimation du courant surfacique océanique par imagerie satellitaire, l'approche permet, dans ses résultats préliminaires, une bonne prise en compte de l'*a priori* parcimonieux avec une complexité linéaire.

4.2 Projet de recherche

Depuis mon arrivée à l'ENSTA Bretagne, mes travaux de recherche s'intègre dans une démarche duale, restant fortement ancrée dans le développement méthodologique et algorithmique, mais s'intéressant à des enjeux défense concrets. Mon projet de recherche approfondit davantage encore cette dualité. A la lecture de mes travaux, de ce qui en fait le cœur, et de mes intérêts particuliers, je distingue deux grands axes de recherche :

- Modèles *a priori* structurés
- Modèles d'observation non-linéaires

Enoncés de la sorte, ces axes de recherche paraissent évidemment très larges et vagues, mais il reflètent bien l'aspect modélisation, central dans mon travail. Contextualisés dans le cadre applicatif de l'acoustique sous-marine passive, ces axes prennent un sens concret intéressant.

Dans chacun de ces grands axes, je souhaite approfondir l'intérêt de l'apprentissage automatique. En revanche, je ne me renierai pas : le cadre bayésien doit me permettre d'apporter des contributions significatives dans chacun d'eux. De fait, pour déployer des algorithmes d'apprentissage automatique, il est connu qu'il faut réunir énormément de données. Or, si on a longtemps considéré comme problématique le foisonnement de données (sous l'appellation de "big data"), nombreuses sont les applications souffrant au contraire d'un manque de données. La problématique se pose particulièrement en acoustique sous-marine passive, où les données recueillies – hors observatoires marins – sont peu nombreuses et les environnements de propagation difficilement reproductibles et transposables. Ainsi, le cadre de l'acoustique sous-marine pose des problèmes existentiels à l'apprentissage automatique, appelant naturellement à conserver un peu d'intelligence "naturelle" aux côtés de l'intelligence artificielle....

4.2.1 Modèles *a priori* structurés

A. Apprentissage automatique pour équations différentielles ordinaires

Soutenue par l'ANR depuis 2020, la chaire IA OceaniX (<https://cia-oceanix.github.io/about/>), dans laquelle l'ENSTA Bretagne est partie prenante, a pour objectif d'établir un lien entre les paradigmes axés sur les modèles qui sous-tendent les sciences physiques et les approches basées sur l'apprentissage axé sur les données, afin de déduire de nouvelles représentations efficaces de systèmes dynamiques complexes, qui fourniront de nouveaux moyens pour la compréhension et le contrôle des océans.

La chaire s'intéresse plus particulièrement aux modèles d'assimilation de données, *i.e.*, à la reconstruction et la prédiction de l'évolution temporelle d'un phénomène physique à partir d'observations satellitaires et/ou *in situ*. Le phénomène est alors supposé répondre à une équation différentielle ordinaire que l'on souhaite apprendre et intégrer dans le problème d'estimation. Plus formellement, si $\mathbf{x}(t)$ est l'état du processus d'intérêt à l'instant t et $\mathbf{y}(t)$ son observation (éventuellement partielle), on aura affaire à un système d'équations du type :

$$\begin{cases} \frac{d}{dt}\mathbf{x}(t) &= \mathcal{M}(\mathbf{x}(t)) \\ \mathbf{y}(t) &= \mathcal{H}_t(\mathbf{x}(t)) + \mathbf{n}(t) \end{cases} \quad (4.1)$$

où \mathcal{M} représente le modèle dynamique, \mathcal{H}_t est l'opérateur d'observation à l'instant t et $\mathbf{n}(t)$ un bruit de mesure. L'apprentissage du modèle décrit par l'équation différentielle ordinaire peut être effectué de plusieurs façons selon différentes architectures neuronales [10]. L'objectif de la chaire est de valider ces modèles au regard des problèmes inverses les intégrant comme *a priori*, mais, au delà également, de proposer des procédures apprenant conjointement modèle et estimateur.

En cours : une thèse En novembre 2021, une thèse – portée par Perrine Bauchot et co-encadrée avec l'IMT Atlantique et le LOPS – a démarré sur la mise en œuvre de stratégies de prises de mesures pour optimiser les performances d'outils de prédiction et reconstruction de l'évolution temporelle de phénomènes physiques. Le sujet s'intéresse ainsi à l'apprentissage conjoint du modèle *a priori* décrivant l'évolution temporelle (selon une équation différentielle ordinaire) et de l'opérateur d'observation, ainsi que de leur bonne prise en compte dans la reconstruction des états du phénomène physique. Au delà des questions portant sur l'apprentissage du modèle physique (choix et implémentation), se pose la question des *a priori* pertinents sur l'opérateur d'observation (parcimonie des mesures prises au regard d'un coût global, trajectoire des mesures, etc.) et de leur prise en compte dans son apprentissage.

Perspectives En acoustique sous-marine, la propagation d'onde est régie par une équation aux dérivées partielles. A ma connaissance, aucune contribution ne s'est encore intéressée à l'apprentissage automatique inférée par cette équation dans ce contexte particulier. Pourtant, des résultats existent qui pourraient être – d'un point de vue purement méthodologique – particularisés à notre cas [29]. Il reste cependant que le milieu marin présente une complexité difficile à appréhender dans une base d'apprentissage exhaustive. Le cas de l'estimation modale en ultra-basses fréquences et milieux petit-fond est une bonne porte d'entrée vers des modèles hybrides basés modèles et apprentissage.

Sur l'optimisation des schémas d'échantillonnage, en revanche, le terrain de l'acoustique sous-marine est plus évident. La problématique se pose en effet dès lors que l'on considère des meutes de drones sous-marins ou le déplacement d'une antenne tractée par exemple. Dans la thèse de Perrine Bauchot, ce sont des techniques issues de l'optimisation et plus spécifiquement, celles classiques de l'apprentissage neuronal qui sont envisagées mais des perspectives ouvrant notamment sur le bayésien experimental design [12] pourraient être considérées par la suite. Ces dernières méthodes pourraient être naturellement couplées avec les approximations variationnelles bayésiennes dans un schéma d'estimation conjointe.

B. Estimation modale en UBF petits fonds

Le post-doctorat de Clément Dorffer sur le sujet a permis de marquer un premier pas vers l'apprentissage de propriétés physiques (relation de dispersion) par le biais de modèles parcimonieux structurés. Cependant, l'apprentissage basé sur la génération de données simulées, *i.e.*, assujetties à des configurations de guides d'ondes et de paramètres environnementaux donnés, n'est pas satisfaisant. Un axe de

recherche que j'aimerais approfondir est constitué par les processus gaussiens [30]. Considérés comme modèles *a priori* sur les amplitudes modales, leurs fonctions de covariance devraient permettre à la fois d'explicitier analytiquement la relation de dispersion liant les nombres d'onde et de modéliser les variations d'amplitudes d'un mode selon les fréquences. Les processus gaussiens ont été largement étudiés dans le cadre de l'apprentissage automatique (on montre notamment qu'ils sont équivalents à des réseaux de neurones infinis [22]). Ils présentent un cadre générique appréciable, donnant à la fois la main sur la structure, la forme, de la fonction de covariance, mais facilitant par ailleurs l'apprentissage de leurs hyperparamètres. A cet égard, certains travaux se sont intéressés à la prise en compte de données d'apprentissage de "basse fidélité" [28] qui me semblent pertinents pour prendre en compte le biais introduit par la simulation (forcément incomplète) des données d'apprentissage. Enfin, je n'oublie pas que ces modèles pourront être avantageusement couplés à des approches par dictionnaires continus et, de par leur caractère gaussien, pris en compte très facilement dans des approximations variationnelles bayésiennes.

4.2.2 Modèles d'observation non-linéaires

A. Imagerie acoustique en mouvement

Dans la continuité directe de mes travaux actuels, les modèles non-linéaires posés par les bruits de phase trouvent un intérêt applicatif dans le problème de localisation de sources (acoustiques) en présence d'incertitudes sur la position des capteurs. En champ proche par exemple, une incertitude sur la position des capteurs aura deux incidences :

- un écart sur le facteur multiplicatif de la direction d'arrivée,
- un écart sur l'atténuation géométrique.

Une manipulation rapide de ces deux grandeurs montre que le premier entrainera potentiellement plus d'erreurs d'estimation sur la direction d'arrivée que le second. Le recours à un modèle de bruit de phase informé pourra alors aider à réaliser la localisation de sources, via son intégration dans la résolution du problème inverse.

Court et moyen termes : une collaboration Ce problème est d'intérêt pour la défense, notamment dans le cadre de la maintenance des bâtiments de surface ou sous-marins. Ainsi une des opérations de contrôle consiste à imager les bâtiments pour en déterminer les réponses acoustiques. Les bâtiments passent parallèlement à l'antenne de capteurs, tractée par un autre navire. Outre l'incertitude autour de la position des capteurs (évoluant dans le temps), la prise en compte de la trajectoire du bâtiment imagé relativement au bâtiment imageur, la présence d'effet Doppler et le faible rapport signal à bruit, posent des écueils supplémentaires. Dans ce contexte, il est primordial de bien poser les hypothèses sur l'environnement et les sources afin d'en tirer profit au maximum. Ainsi, on pourra considérer un modèle bayésien prenant en compte un bruit de phase sur chaque capteur en informant les phases non seulement par les liens inter-capteurs (comme envisagé notamment dans la thèse de Guillaume Beaumont sur les fluctuations océaniques) mais également par les liens inter-snapshots (*i.e.* entre deux échantillons consécutifs reçus sur l'antenne). Ce modèle sera couplé à un tracking des sources dont le glissement de fréquences dû à l'effet Doppler doit être filé d'un snapshot à l'autre.

Suite à des discussions avec le centre de Brest de DGA Techniques Navales, une première étude (achevée en mars 2021) nous a permis d'établir les lignes directrices de ce projet. Une thèse doit démarrer à la rentrée 2022.

Perspectives Au delà de l'intérêt scientifique, je vois dans cet axe de recherche l'opportunité et l'importance de renforcer des liens entre des entités locales, brestoises. Ainsi, ce projet implique l'IMT Atlantique mais intéresse plus particulièrement le centre de Brest de DGA Techniques Navales qui sera également partie prenante de l'encadrement de la thèse. L'apport de chacun y est également apprécié et devrait, je l'espère, porter des fruits sur un terme plus long que le seul cadre de la thèse.

B. Fluctuations océaniques

Mes travaux sur les traitements d'antenne en milieu fluctuant et ceux de Guillaume Beaumont qui a poursuivi en thèse avec moi, ont permis de proposer des solutions bayésiennes innovantes, générant néanmoins la frustration de ne pouvoir lier leurs paramètres aux variables physiques décrivant les fluctuations océaniques. De fait, les recherches sur les fluctuations océaniques – la connaissance des phénomènes océanographiques et leurs impacts sur la propagation acoustique – et celles portant sur les traitements d'antenne – détection et localisation de sources – ont évolué de façon parallèle, ignorant les avancées voisines et empêchant leur transfert de l'une vers l'autre. Or, si l'on souhaite dépasser le stade de l'utilisation d'outils sur étagères de part et d'autre, ce transfert est maintenant nécessaire.

En cours : une thèse En janvier 2021, une thèse CIFRE Défense – portée par Alexandre L'Her – a démarré sur ce sujet en collaboration avec le Shom, DGA Techniques Navales et Thales Underwater Systems. La thèse se propose d'étudier, formaliser et exploiter les liens entre phénomènes océanographiques aléatoires, propagation acoustique en milieu fluctuant et traitements d'antenne adaptatifs. Pour ce faire, trois grands volets sont abordés : i) compréhension et quantification des phénomènes océanographiques à l'origine des fluctuations de l'environnement, ii) étude et simulation de la propagation acoustique dans des environnements fluctuants, iii) validation d'outils de traitement de signal acoustique adaptés sur données simulées et données acquises en milieu réel. La finalité est de faire le lien entre l'océanographie et l'acoustique sous-marine, afin d'alimenter les *a priori* des algorithmes de localisation robustes.

La thèse se veut plus englobante que celle de Guillaume Beaumont dans laquelle nous posions comme pré-requis l'hypothèse d'un bruit de phase le long de l'antenne. Cette hypothèse restreint en effet l'étude à des régimes de fluctuations particuliers dans lesquels l'onde acoustique est supposée traverser des masses d'eau identiques quelle que soit sa direction d'arrivée. De plus, elle adopte un point de vue "antenne" dont il peut être difficile de s'affranchir dès lors que l'on souhaite l'informer par des paramètres environnementaux.

La piste explorée par Alexandre L'Her est celle de l'apprentissage de modèles liant les paramètres environnementaux relevés *in situ* (thermistance) aux observations acoustiques. Parmi les outils possibles, celui de l'analyse des corrélations des deux ensembles de variables – acoustiques et environnementales – (CCA pour canonical correlation analysis) offre des perspectives intéressantes. Ce type d'approche doit constituer une façon simple de lier explicitement les paramètres environnementaux (ex. température) à ceux utiles à la détection et estimation de directions d'arrivée (ex. rayon de cohérence de l'antenne) par un modèle linéaire. Cette relation doit *in fine* alimenter un modèle couplant les linéarités attendues d'une propagation dans un milieu homogène et stationnaire et les non-linéarités aléatoires induites par les fluctuations du milieu.

Perspectives La thèse d'Alexandre L'Her est ambitieuse, pluri-disciplinaire et prospective. D'un point de vue purement signal – celui qui m'occupe – elle ouvre des voies qui n'auront certainement pas toutes trouvé leurs réponses dans deux ans. Ainsi, le couplage apprentissage/modèle posé par le problème peut être envisagé de différentes manières : accentuant l'aspect apprentissage, d'autres techniques prenant

en compte des non-linéarités (à l'instar de la Kernel CCA [35] par exemple) pourraient envisagées ; privilégiant le modèle, on pourrait au contraire guider davantage l'apprentissage par sa restriction à certaines variables latentes d'intérêt. A noter que le choix même d'une pondération plutôt qu'une autre (modèle ou apprentissage) dépend très fortement des cas d'études et des régimes de fluctuations considérés.

4.3 Synthèse et ouverture

Jusqu'à présent, mes recherches m'ont conduite à développer de nouvelles méthodes de traitement du signal prenant en compte des modèles physiques. En parallèle, j'ai pu assister à l'avènement des techniques d'apprentissage automatique, dont on ne peut nier les grandes performances dès lors que des données sont disponibles en grandes quantité et diversité. Couplant ces deux aspects et au delà des axes de recherche identifiés ci-dessus, je souhaite faire des prochaines années un terrain d'investigation pour des méthodes prenant le meilleur des deux mondes :

- L'apprentissage, contraint par des modèles physiques, doit permettre de recourir à moins de données mais aussi, *a contrario*, d'accéder à des informations sur la physique sous-tendant les phénomènes mesurés. Ainsi, on peut imaginer passer de la localisation de sources *malgré* les fluctuations océaniques à la localisation de sources *et* la caractérisation des fluctuations océaniques, voire même comme dans le cas des milieux diffusants, considérer les fluctuations comme une source de richesse/ "diversité" - au sens signal - pouvant peut-être être exploitées à leur avantage (à l'image des problèmes inverses randomisés [15]).
- Enfin, l'apprentissage, couplé à l'inférence bayésienne, doit permettre de conserver la quantification des (in)certitudes permise par le cadre probabiliste, tout en tirant profit de la puissance de modélisation des architectures neuronales. L'approche est particulièrement d'intérêt dès lors que l'on se confronte aux erreurs de modèles, notamment rencontrées dans les milieux fluctuants mais également en imagerie acoustique par antenne tractée.

Bibliographie

- [1] M. Bayati and A. Montanari. The dynamics of message passing on dense graphs, with applications to compressed sensings. *IEEE Transactions on Information Theory*, 57(2), 2011.
- [2] G. Beaumont. *Traitement correctif des effets de décohérence acoustique induits par les fluctuations du milieu de propagation : algorithmes d'estimation de directions d'arrivée en milieu fluctuant*. PhD thesis, ENSTA Bretagne, Juillet 2020.
- [3] F. Caltagirone, L. Zdeborová, and F. Krzakala. On convergence of approximate message passing. *IEEE International Symposium on Information Theory*, 2014.
- [4] E. J. Candès, T. Strohmer, and V. Voroninski. Phaselift : exact and stable signal recovery from magnitude measurements via convex programming. *Communications in Pure and Applied Mathematics*, 66(8), 2013.
- [5] E. J. Candès, M. B. Wakin, and S. Boyd. Enhancing sparsity by reweighted l1 minimization. *Journal of Fourier Analysis and Applications*, 14(5), 2008.
- [6] M. Cannon, C. Cullum, and E. Polak. *Theory of Optimal Control and Mathematical Programming*. New York, 1970.
- [7] S. S. Chen and D. L. Donoho. Basis pursuit. In *Proc. Asilomar Conference on Signals, Systems, and Computers*, volume 1, pages 41–44, 1994.
- [8] J. A. Colosi and M. G. Brown. Efficient numerical simulation of stochastic internal-wave induced sound-speed perturbation fields. *Journal of Acoustical Society of America (JASA)*, 103(4), 1998.
- [9] D. L. Donoho, A. Javanmard, and A. Montanari. Information- theoretically optimal compressed sensing via spatial coupling and approximate message passing. *IEEE transactions on information theory*, 59(11), 2013.
- [10] R. Fablet, B. Chapron, L. Drumetz, E. Mémin, O. Pannekoucke, and F. Rousseau. Learning variational data assimilation models and solvers. *Journal of Advances in Modeling Earth Systems*, 13, 2021.
- [11] J. R. Fienup. Phase retrieval algorithms : a comparison. *Applied Optics*, 21(15), 1982.
- [12] A. Foster, M. Jankowiak, E. Bingham, P. Horsfall, Y. W. Teh, T. Rainforth, and N. Goodman. Variational bayesian optimal experimental design. *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [13] S. Foucart and H. Rauhut. *A mathematical introduction to compressive sensing*. Birkhaeuser, 2010.
- [14] R. Gerchberg and W. Saxton. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik*, 35, 1972.
- [15] R. M. Gower. *Sketch and Project : Randomized Iterative Methods for Linear Systems and Inverting Matrices*. PhD thesis, The University of Edinburgh, 2016.

- [16] R. Gribonval and M. Nielsen. Sparse representations in unions of bases. *IEEE Trans. On Information Theory*, 49(12) :3320–3325, December 2003.
- [17] D. Griffin and J. Lim. Signal estimation from modified short-time fourier transform. *IEEE Transactions On Acoustics, Speech and Signal Processing*, 32(2), 1984.
- [18] G. E. Hinton. A practical guide to training restricted boltzmann machines. *Neural Networks : Tricks of the trade*, 7700 :599–619, 2012.
- [19] F. B. Jensen, W. A. Kuperman, M. B. Porter, and H. Schmidt. *Computational ocean acoustics*. 2011.
- [20] F. Krzakala, A. Manoel, E. W. Tramel, and L. Zdeborová. Variational free energies for compressed sensing. In *Proc. IEEE Int. Symposium on Information Theory*, pages 1499–1503, June 2014.
- [21] F. Krzakala, M. Mézard, F. Sausset, Y. F. Sun, and L. Zdeborová. Statistical-physics-based reconstruction in compressed sensing. *Physical Review X*, 2 :021005, May 2012.
- [22] J. Lee, Y. Bahri, R. Novak, S. S. Schoenholz, J. Pennington, and J. Sohl-Dickstein. Deep neural networks as gaussian processes. *Proc. Int. Conf. on Learning Representations*, 2018.
- [23] X. Li and V. Voroninski. Sparse signal recovery from quadratic measurements via convex programming. *SIAM Journal of Mathematical Analysis*, 45(5), 2013.
- [24] S. Mallat. *A Wavelet Tour of Signal Processing - The Sparse Way*. Academic Press, third edition, December 2008.
- [25] S. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *IEEE Trans. On Signal Processing*, 41(12) :3397–3415, December 1993.
- [26] H. Ohlsson, A. Y. Yang, R. Dong, and S. S. Sastry. Compressive phase retrieval from squared output measurements via semidefinite programming. *IFAC Symposium on System Identification*, 16(1), 2012.
- [27] Y. C. Pati, R. Rezaifar, and P. S. Krishnaprasad. Orthogonal matching pursuit : Recursive function approximation with applications to wavelet decomposition. In *Proc. Asilomar Conference on Signals, Systems, and Computers*, pages 40–44, 1993.
- [28] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Inferring solutions of differential equations using noisy multi-fidelity data. *Journal of Computational Physics*, 335 :736–746, 2017.
- [29] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks : A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378(7), 2019.
- [30] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. 2005.
- [31] G. Real. *An ultrasonic testbench for reproducing the degradation of sonar performance in a fluctuating ocean*. PhD thesis, Aix-Marseille Université, 2015.
- [32] N. Le Roux and Y. Bengio. Representational power of restricted boltzmann machines and deep belief networks. *Neural computation*, 20(6), 2008.
- [33] Y. Schechtman, A. Beck, and Y. C. Eldar. Gespar : Efficient phase retrieval of sparse signals. *IEEE Transactions on Signal Processing*, 62(4), 2013.
- [34] P. Schniter and S. Rangan. Compressive phase retrieval via generalized approximate message passing. *Communication, Control, and Computing (Allerton)*, 2012.
- [35] A. Tenenhaus, C. Philippe, and V. Frouin. Kernel generalized canonical correlation analysis. *Computational Statistics & Data Analysis*,, 2015.

- [36] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society*, 58 :267–288, 1996.
- [37] I. M. Vellekoop and A. P. Mosk. Focusing coherent light through opaque strongly scattering media. *Optics Letters*, 32(16), 2007.
- [38] J. Vila and P. Schniter. Expectation-maximization bernoulli-gaussian approximate message passing. In *Proc. Asilomar Conf. on Signals, Systems, and Computers*, November 2011.
- [39] I. Waldspurger, A. d’Aspremont, and S. Mallat. Phase recovery, maxcut and complex semidefinite programming. *Mathematical Programming Series A- Springer*, 2013.
- [40] Y. Wu and S. Verdu. Optimal phase transitions in compressed sensing. *IEEE Transactions on Information Theory*, 58, 2012.
- [41] A. Xenaki and P. Gerstoft. Grid-free compressive beamforming. *JASA*, 137(4), 2015.
- [42] J. S. Yedidia, W. T. Freeman, and Y. Weiss. Constructing free-energy approximations and generalized belief propagation algorithms. *IEEE Transactions on Information Theory*, 51(7) :2282–2312, July 2005.

Annexe A

Curriculum Vitae

Angélique Drémeau

née le 1^{er} juillet 1982 à Enghien-les-bains (95)

<i>Adresse</i>	Lab-STICC & ENSTA Bretagne 2 rue François Verny, 29806 Brest Cedex 9
<i>Téléphone</i>	02 98 34 89 71
<i>Courriel</i>	angelique.dremeau@ensta-bretagne.fr
<i>Page personnelle</i>	http://angelique.dremeau.free.fr/

Situation actuelle

Depuis Jan. 2015 **Enseignante-Chercheuse**, Ecole Nationale Supérieure de Techniques Avancées Bretagne (ENSTA Bretagne), Lab-STICC (UMR 6285), DMID/MATRIX, Brest.
(En congé maternité de septembre 2018 à février 2019)

Thématiques de recherche

<i>Compétences</i>	Outils statistiques : modélisation et inférence bayésiennes, approximations variationnelles bayésiennes, méthodes d'échantillonnage. Représentations parcimonieuses et structurées : algorithmes de décomposition, apprentissage de dictionnaire.
<i>Contributions</i>	Problèmes inverses : modélisation, algorithmie, notamment en reconstruction de phase (informée) et dictionnaires continus. Acoustique sous-marine passive : estimation de directions d'arrivée, estimation modale.

Formation

<i>2010</i>	Doctorat , mention Traitement du Signal et Télécommunications, Université de Rennes 1, distinction Très honorable.
<i>2007</i>	Master Recherche "Signal, Télécommunications, Image", Université de Rennes 1, mention Bien.
<i>2007</i>	Diplôme d'ingénieur de l'Ecole Nationale Supérieure des Télécommunications de Bretagne (maintenant IMT Atlantique), Brest.

Recherche

- Oct. 2013 -
Dec. 2014
- Post-doctorat**, Ecole Normale Supérieure (ENS), Paris
Sujet : Problèmes inverses en imagerie sous contrainte de parcimonie
Direction : Florent Krzakala
- Sept. 2011 -
Sept. 2013
- Post-doctorat**, Télécom ParisTech, Paris
Sujet : Analyse de mouvement par acquisition multimodale
Direction : Gaël Richard
- Dec. 2010 -
Août 2011
- Post-doctorat**, Institut Langevin, Paris
Sujet : Représentation minimale et intelligibilité des sons
Direction : Laurent Daudet
- Oct. 2007 -
Nov. 2010
- Doctorat**, INRIA, IRISA, Rennes
Sujet : Décompositions parcimonieuses : approches bayésiennes et application à la compression d'image.
Soutenance : le 19/11/2010 devant le jury composé de Laurent Daudet (Président), Pierre Vanderghyest, Jérôme Idier (Rapporteurs), Béatrice Pesquet-Popescu (Examinatrice), Jean-Jacques Fuchs, Christine Guillemot (Co-directeurs de thèse).

Projets

- Août 2020 -
Août 2025
- OceaniX**, Physics-Informed AI for Observation-driven Ocean AnalytiX
Porteur : Ronan Fablet (IMT Atlantique)
Statut : chaire ANR en Intelligence Artificielle, 2M€
Rôle : membre du consortium
- Oct. 2019 -
Oct. 2021
- Ph-IA**, Physics-guided probabilistic deep learning for underwater acoustics
Porteur : Gilles Le Chenadec (ENSTA Bretagne)
Statut : projet AID, 131k€
Rôle : partenaire
- Oct. 2015 -
Oct. 2018
- TS-DECO**, Traitements correctifs des effets de décohérence acoustique induits par les fluctuations du milieu de propagation
Statut : projet AID, 197k€
Rôle : porteuse du projet
- Sept. 2011 -
Sept. 2013
- 3DLIFE**, Bringing the Media Internet to Life
Porteur : Queen Mary University of London
Statut : réseau d'excellence européen FP7, 4M€
Rôle : coordinatrice de workpackage
- REVERIE**, REal and Virtual Engagement in Realistic Immersive Environments
Porteur : ST Microelectronics
Statut : projet européen collaboratif FP7, 10M€
Rôle : coordinatrice de workpackage

Expertise et rayonnement scientifique

- Publications* – 12 articles dans des journaux internationaux à comités de lecture dont 9 dans des journaux de facteurs d'impact supérieurs à 2 (IEEE Trans. on Signal Processing, IEEE Journal of Oceanic engineering...).
- 27 articles dans des actes de conférences internationales à comités de lecture.
- Séminaires invités* – mai 2017, au Cercle Naval de Brest dans le cadre de "l'atelier acoustique sous-marine" organisé par la start-up SIVIENN
- janvier 2017, au Laboratoire d'Imagerie Biomédicale, Jean-Gabriel Minonzio
- mars 2015, à l'institut Langevin, Laurent Daudet
- Collaborations* – collaboration avec Florent Le Courtois (Shom) et Gaultier Real (DGA TN) sur la thématique de propagation acoustique en milieu océanique fluctuant
- collaboration avec Julien Bonnel (WHOI), avec Barbara Nicolas (CNRS, Creatis) et Thierry Chonavel (IMT Atlantique) sur la thématique de reconstruction de modes en acoustique sous-marine UBF
- collaboration avec Ronan Fablet (IMT Atlantique) et Gilles Le Chenadec (ENSTA Bretagne) sur la thématique de l'intelligence artificielle inspirée par la physique
- collaboration avec Cédric Herzet (INRIA Rennes) sur la thématique des représentations parcimonieuses, notamment sur dictionnaires continus.
- Examinatrice dans des jurys de soutenance de thèses* – Van Duong Nguyen, 12/2020, IMT Atlantique
direction : Ronan Fablet et René Garello
- Maëlle Torterotot, 11/2020, Université de Bretagne Occidentale
direction : Jean-Yves Royer et Flore Samaran
- Mathieu Fontaine, 06/2019, Université de Lorraine
direction : Roland Badeau et Antoine Liutkus
- Clément Dorffer, 12/2017, Université du Littoral Côte d'Opale
direction : Gilles Roussel et Matthieu Puigt
- Mohammad Ammad-Uddin, 12/2017, Université Bretagne Loire
direction : Ali Mansour
- Membre de comité d'organisation* SERENADE, workshop sur la Surveillance, Etude et Reconnaissance de l'Environnement mariN par Acoustique Discrète
- évènement biennal national structurant la communauté Acoustique Passive et Discrète, \simeq 80 participants.
- Reviewer* pour les revues IEEE Trans. on Signal Processing, Elsevier Signal Processing, IEEE Journal of Oceanic Engineering, The Journal of Acoustical Society of America, pour les conférences GRETSI, ICASSP, iTWIST, SPARS...
- Qualification* aux fonctions de maître de conférences, section 61, 2012.

Enseignement

* Responsabilités pédagogiques

Depuis Jan. 2015 **Traitement du signal**
 Dispensé en première année du cycle ingénieur (CI), ce cours fait partie du tronc commun et concerne l'ensemble de la promotion des étudiants (\simeq 180 étudiants). J'en assure les cours magistraux, l'examen, une partie des TD et TP, et suis en charge de la planification des interventions de mes collègues en tant qu'encadrants de TD/TP.

Depuis Jan. 2020 **Traitement du signal avancé**
 Dispensé en deuxième année du cycle ingénieur, ce cours fait partie de la spécialité SOIA (systèmes d'observation et intelligence artificielle, \simeq 25 étudiants). J'en assure les cours magistraux, l'examen et les TD/TP.

Représentations parcimonieuses et acquisition compressée

Dispensé en deuxième année du cycle ingénieur, ce cours est proposé en option de la spécialité SOIA (systèmes d'observation et intelligence artificielle, \simeq 25 étudiants). J'en assure les cours magistraux, l'examen et les TD/TP.

* Charge d'enseignement pour l'année universitaire 2019-2020 à l'ENSTA Bretagne

(NB : ne sont mentionnés que les modules dont le volume horaire est supérieur à 10h équivalent TD)

<i>1ère année CI</i>	– Cours/TD/TP Traitement du signal (46 h)
	– TD/TP Probabilités et statistiques (21 h)
<i>2ème année CI</i>	– Cours/TD/TP Traitement du signal avancé (20h)
	– Cours/TD/TP Représentations parcimonieuses, acquisition compressée (24h)
<i>2ème année CI par alternance</i>	– Cours/TD/TP Traitement du signal (22h) (à destination d'élèves en formation d'ingénieur par alternance en entreprise)

Ma charge annuelle (\simeq 180h équivalent TD) est complétée, entre autres, par des interventions en apprentissage automatique (2ème année CI), automatique (1ère année CI) ou en formation continue (acoustique sous-marine).

* Pendant mon doctorat et mes post-doctorats

*Oct. 2007 -
Dec. 2014* Vacations régulières en méthodes d'optimisation, traitement du signal et de l'image, et encadrement de projets à l'université de Rennes 1 & 2, l'ISEP et Telecom ParisTech pour un volume global d'environ 110 heures.

Encadrement

Stage de Master 2 **Alexandre L'Her** (avril - août 2020)
Sujet : The effect of internal waves on sound propagation : a case study of the Messina strait.
Co-encadrement : Florent Le Courtois (Shom), Gaultier Real (DGA TN).
 A. L'Her a débuté une thèse de doctorat en janvier 2021 sous notre encadrement.

Fabio Cassiano (février - août 2017)
Sujet : Détection de clics d'écholocation d'orques et de cachalots
Co-encadrement : Flore Samaran (ENSTA Bretagne)
 F. Cassiano a débuté une thèse sous notre encadrement.

Sabrina Bourmani (avril - août 2016)
Sujet : Estimation de directions d'arrivée en milieu océanique fluctuant
 S. Bourmani a poursuivi et soutenu une thèse de doctorat sous le co-encadrement de D. Pastor et F.-X. Socheleau (IMT Atlantique).

Doctorat **Alexandre L'Her** (janvier 2021 -)
Sujet : Propagation acoustique en milieu fluctuant : des phénomènes océanographiques au traitement du signal.
Financement : CIFRE Défense (Thalès)
Direction : Yann Stephan (Shom)
Co-encadrement : Florent Le Courtois (Shom), Gaultier Real (DGA TN).
Taux d'encadrement : 25%

Fabio Cassiano (avril 2019 -)
Sujet : Détection de clics d'écholocation d'orques et de cachalots
Financement : CIFRE (Fondation d'entreprises des terres australes)
Direction : Isabelle Quidu (ENSTA Bretagne)
Co-encadrement : Flore Samaran (ENSTA Bretagne)
Taux d'encadrement : 25%
 La thèse de F. Cassiano a malheureusement été interrompue.

Thomas Paviet-Salomon (octobre 2017 - mai 2021)
Sujet : Super-résolution modale à partir d'une antenne linéaire horizontale
Financement : bourse DGA et ONRG
Direction : Thierry Chonavel (IMT Atlantique)
Co-encadrement : Barbara Nicolas (CNRS, Creatis)
Taux d'encadrement : 50%
 T. Paviet-Salomon est actuellement en post-doctorat.

Guillaume Beaumont (octobre 2016 - juillet 2020)
Sujet : Algorithmes d'estimation de directions d'arrivée en milieu fluctuant
Financement : contrat DGA
Direction : Ronan Fablet (IMT Atlantique)
Taux d'encadrement : 50%
 G. Beaumont a intégré DGA Techniques Navales en tant qu'expert en acoustique sous-marine à l'automne 2021.

Post-doctorat

Clément Dorffer (janvier 2018 - mars 2022)

Sujet : Développement d'outils d'intelligence artificielle pour l'estimation modale en acoustique sous-marine.

Financement : projet AID

Co-encadrement : Gilles Le Chenadec (ENSTA Bretagne).

Riwal Lefort (décembre 2015 - juin 2017)

Sujet : Traitement d'antenne en milieu océanique fluctuant

Financement : projet AID

Administration et responsabilités collectives

Depuis Avril 2021

Représentante du Lab-STICC sur le site de l'ENSTA Bretagne.

Jan. 2017 - Mars 2022

Représentante élue au conseil de laboratoire du Lab-STICC.

Sep. 2017 - Déc. 2020

Représentante de l'équipe Lab-STICC/TOMS à l'ENSTA Bretagne.

Annexe B

Publications et communications

Toutes les publications mentionnées ci-dessous sont disponibles sur ma page personnelle
<http://angelique.dremeau.free.fr>

Chapitre de livre

- [L] A. Drémeau, C. Schülke, Y. Xu, D. Shah – **Statistical inference with probabilistic graphical models** – Dans “*Statistical Physics, Optimization, Inference, and Message-Passing Algorithms*”, Eds : F. Krzakala, F. Ricci-Tersenghi, L. Zdeborova, R. Zecchina, E. W. Tramel, L. F. Cugliandolo (Oxford University Press), arXiv :1409.4928, 2014.

Revue internationale

- [R12] T. Paviet-Salomon, J. Bonnel, C. Dorffer, B. Nicolas, T. Chonavel, D. Tollefsen, D. P. Knobles, P. S. Wilson, A. Drémeau – **Estimation of frequency-wavenumber diagrams using a physics-based grid-free compressed sensing method** – *IEEE Journal of Oceanic Engineering*, accessible en ligne, Octobre 2021.
- [R11] C. Herzet, C. Dorffer, A. Drémeau – **Gather and Conquer : Region-based Strategies to Accelerate Safe Screening Tests** – *IEEE Transactions On Signal Processing*, Vol. 67, No. 12, pp. 3300 - 3315, Mai 2019.
- [R10] C. Herzet, P. Héas, A. Drémeau – **Model reduction from partial observations** – *Int’l Journal for Numerical Methods in Engineering*, Vol.113, No. 3, pp. 479-511, Janvier 2018.
- [R9] R. Lefort, R. Emmetière, S. Bourmani, G. Real, A. Drémeau – **Sub-antenna processing for coherence loss in underwater direction-of-arrival estimation** – *The Journal of the Acoustical Society of America (JASA)*, Vol.142, No. 4, pp. 2143-2154, 2017.
- [R8] A. Drémeau, F. Le Courtois, J. Bonnel – **Reconstruction of dispersion curves in the frequency-wavenumber domain using compressed sensing on a random array** – *IEEE Journal of Oceanic Engineering*, Vol. 42, No. 4, pp. 914 - 922, Mars 2017.
- [R7] R. Lefort, G. Real, A. Drémeau – **Direct regressions for underwater acoustic source localization in fluctuating oceans** – *Elsevier Applied Acoustics*, Vol. 116, pp. 303-310, Janvier 2017.

- [R6] E. W. Tramel, A. Drémeau, F. Krzakala – **Approximate Message Passing with Restricted Boltzmann Machine Priors** – *Journal of Statistical Mechanics : Theory and Experiment*, Vol. 2016, No. 7, pp. 073401, 2016.
- [R5] P. Héas, A. Drémeau, C. Herzet – **An efficient algorithm for video superresolution based on a sequential model** – *SIAM Journal on Imaging Sciences*, Vol. 9, No. 2, pp. 537-572, 2016.
- [R4] C. Herzet, A. Drémeau, C. Soussen – **Relaxed Recovery Conditions for OMP/OLS by Exploiting both Coherence and Decay** – *IEEE Transactions On Information Theory*, Vol. 62, No. 1, pp. 1-12, Janvier 2016.
- [R3] A. Drémeau, A. Liutkus, D. Martina, O. Katz, C. Schülke, F. Krzakala, S. Gigan, L. Daudet – **Reference-less measurement of the transmission matrix of a highly scattering material using a DMD and phase retrieval techniques** – *Optics Express*, Vol. 23, pp. 11898-11911, 2015.
- [R2] C. Sued, A. Drémeau, D. Pressnitzer, L. Daudet – **Auditory sketches : Sparse representations of sounds based on perceptual models** – *Lecture Notes in Computer Science (LNCS)*, Springer Verlag, Vol. 7900, pp. 154-170, 2013.
- [R1] A. Drémeau, C. Herzet, L. Daudet – **Boltzmann machine and mean-field approximation for structured sparse decompositions** – *IEEE Transactions On Signal Processing*, Vol. 60, No. 7, pp. 3425 - 3438, Juillet 2012.

Actes de conférences internationales

- [A27] C. Dorffer, T. Paviet-Salomon, G. Le Chenadec, A. Drémeau – **Modal estimation in underwater acoustics by data-driven structures sparse decomposition** – *European Signal Processing Conference (EUSIPCO)*, conférence virtuelle, Septembre 2021.
- [A26] T. Paviet-Salomon, C. Dorffer, J. Bonnel, B. Nicolas, T. Chonavel, A. Drémeau – **Dispersive grid-free orthogonal matching pursuit for modal estimation in ocean acoustics** – *IEEE Int'l Conference on Acoustics, Speech and Signal Processing (ICASSP)*, conférence virtuelle, Mai 2020.
- [A25] C. Dorffer, C. Herzet, A. Drémeau – **A Region-based relaxations to accelerate greedy approaches** – *European Signal Processing Conference (EUSIPCO)*, Coruna, Espagne, Septembre 2019.
- [A24] G. Beaumont, A. Drémeau, R. Fablet, G. Real, F. le Courtois – **At-sea experiment evaluation of the influence of environmental fluctuations on the acoustic coherence radius** – *Int'l Conference and Exhibition on Underwater Acoustics (UACE)*, Hersonissos, Crete, Grèce, Juillet 2019.
- [A23] G. Beaumont, A. Drémeau, R. Fablet – **Von Mises prior for phase-noisy DOA estimation : the VITAMIN algorithm** – *Int'l Conference and Exhibition on Underwater Acoustics (UACE)*, Hersonissos, Crete, Grèce, Juillet 2019.
- [A22] G. Beaumont, R. Fablet, A. Drémeau – **An Approximate Message Passing approach for DOA estimation in phase noisy environments** – *Int'l Conference on Latent Variable Analysis and Signal Separation (LVA/ICA)*, Guildford, UK, Juillet 2018.

- [A21] C. Herzet, A. Drémeau – **Joint screening tests for LASSO** – *IEEE Int'l Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Calgary, Canada, Avril 2018.
- [A20] R. Lefort, A. Drémeau – **Sub-Antenna Sparse Processing For Coherence Loss In Underwater Source Localization** – *European Signal Processing Conference (EUSIPCO)*, Kos Island, Grèce, Septembre 2017.
- [A19] A. Drémeau, C. Herzet – **DOA estimation in structured phase-noisy environments** – *IEEE Int'l Conference on Acoustics, Speech and Signal Processing (ICASSP)*, New Orleans, USA, Mars 2017.
- [A18] A. Drémeau, A. Deleforge – **Phase retrieval with a multivariate Von Mises prior : from a Bayesian formulation to a lifting solution** – *IEEE Int'l Conference on Acoustics, Speech and Signal Processing (ICASSP)*, New Orleans, USA, Mars 2017.
- [A17] R. Lefort, A. Drémeau – **Spatio-temporal weighting for underwater source localization in the context of coherence loss** – *Conference on Acoustic and Environmental Variability, Fluctuations and Coherence*, Cambridge, UK, Decembre 2016.
- [A16] A. Saade, F. Caltagirone, I. Carron, L. Daudet, A. Drémeau, S. Gigan, F. Krzakala – **Random Projections through multiple optical scattering : Approximating kernels at the speed of light** – *IEEE Int'l Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Shanghai, Chine, Mars 2016.
- [A15] A. Drémeau, F. Krzakala – **Phase recovery from a Bayesian point of view : the variational approach** – *IEEE Int'l Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Brisbane, Australie, Avril 2015.
- [A14] A. Drémeau, P. Héas, C. Herzet – **Sparse representations in nested non-linear models** – *IEEE Int'l Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Florence, Italie, Mai 2014.
- [A13] F. Rigaud, A. Drémeau, B. David, L. Daudet – **A probabilistic line spectrum model for musical instrument sounds and its application to piano tuning estimation** – *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, Mohonk Mountain House, New Paltz, New York, USA, Octobre 2013.
- [A12] A. Drémeau, S. Essid – **Probabilistic dance performance alignment by fusion of multimodal features** – *IEEE Int'l Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Vancouver, Canada, Mai 2013.
- [A11] R. Badeau, A. Drémeau – **Variational Bayesian EM algorithm for modeling mixtures of non-stationary signals in the time-frequency domain (HR-NMF)** – *IEEE Int'l Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Vancouver, Canada, Mai 2013.
- [A10] A. Liutkus, A. Drémeau, D. Alexiadis, S. Essid, P. Daras – **Analysis of dance movements using Gaussian processes** – *ACM Multimedia Grand Challenge*, Nara, Japon, Octobre 2012.
- [A9] S. Essid, D. Alexiadis, R. Tournemenne, M. Gowing, P. Kelly, D. Monaghan, P. Daras, A. Drémeau, N. E. O'Connor – **An advanced virtual dance performance evaluator** – *IEEE Int'l Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Kyoto, Japon, Mars 2012.
- [A8] A. Drémeau, C. Herzet, L. Daudet – **Structured Bayesian orthogonal matching pursuit** – *IEEE Int'l Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Kyoto, Japon,

Mars 2012.

- [A7] A. Drémeau, C. Herzet, L. Daudet – **Soft Bayesian pursuit algorithm for sparse representations** – *IEEE Int'l Workshop on Statistical Signal Processing (SSP)*, Nice, France, Juin 2011.
- [A6] A. Drémeau, M. Türkan, C. Herzet, C. Guillemot, J.-J. Fuchs – **Spatial intra-prediction based on mixtures of sparse representations** – *IEEE Int'l Workshop on Multimedia Signal Processing (MMSP)*, Saint-Malo, France, Octobre 2010.
- [A5] C. Herzet, A. Drémeau – **Bayesian Pursuit Algorithms** – *European Signal Processing Conference (EUSIPCO)*, Aalborg, Danemark, Août 2010.
- [A4] A. Drémeau, C. Herzet – **An EM-Algorithm Approach for the Design of Orthonormal Bases adapted to Sparse Representations** – *IEEE Int'l Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Dallas, USA, Mars 2010.
- [A3] C. Herzet, A. Drémeau – **Sparse Representation Algorithms based on Mean-Field Approximations** – *IEEE Int'l Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Dallas, USA, Mars 2010.
- [A2] A. Drémeau, C. Herzet, C. Guillemot, J.-J. Fuchs – **Sparse Optimization with Directional DCT Bases for Image Compression** – *IEEE Int'l Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Dallas, USA, Mars 2010.
- [A1] A. Drémeau, C. Herzet, C. Guillemot, J.-J. Fuchs – **Anisotropic Multi-scale Sparse Learned Bases for Image Compression** – *IS&T/SPIE Electronic Imaging*, vol. 7543, San Jose, USA, Janvier 2010.

Actes de conférences nationales

- [N3] G. Beaumont, R. Fablet, A. Drémeau – **Approche de type message passing pour la localisation de source en milieu fluctuant** – *GRETSI*, Lille, France, Août 2019.
- [N2] A. Drémeau, F. Le Courtois, J. Bonnel – **Reconstruction des modes de propagation en acoustique sous-marine par acquisition compressée** – *GRETSI*, Lyon, France, Septembre 2015.
- [N1] A. Drémeau, A. Liutkus, D. Martina, O. Katz, C. Schülke, F. Krzakala, S. Gigan, L. Daudet – **Approches bayésiennes pour la reconstruction de phase – Application à l'optique des milieux complexes** – *GRETSI*, Lyon, France, Septembre 2015.

Communications à des conférences internationales et nationales sans acte

- [C7] C. Dorffer, T. Paviet-Salomon, G. Le Chenadec and A. Drémeau – **Learning sparse structures for physics-inspired compressed sensing** – *Int'l Travelling Workshop on Interactions between Sparse Models and Technology (iTWIST)*, conférence virtuelle, Décembre 2020.
- [C6] T. Paviet-Salomon, C. Dorffer, J. Bonnel, B. Nicolas, T. Chonavel, A. Drémeau – **Dispersive grid-free orthogonal matching pursuit for modal estimation in ocean acoustics** – *Forum Acousticum*, conférence virtuelle, Novembre 2020.

- [C5] A. Drémeau, C. Herzet – **DOA estimation in fluctuating oceans : put your glasses on!** – *Workshop on Signal Processing with Adaptive Sparse Structured Representations (SPARS)*, Lisbon, Portugal, July 2017.
- [C4] A. Drémeau, F. Le Courtois, J. Bonnel – **Estimation des nombres d'onde par acquisition compressée en propagation acoustique sous-marine** – *Conférence française d'acoustique (CFA)*, Le Mans, France, Avril 2016.
- [C3] A. Drémeau, F. Krzakala – **Structured sparsity : towards a « deep » understanding** – *Workshop on Signal Processing with Adaptive Sparse Structured Representations (SPARS)*, Cambridge, UK, Juillet 2015.
- [C2] A. Drémeau, P. Héas, C. Herzet – **Combining sparsity and dynamics : an efficient way** – *Int'l Travelling Workshop on Interactions between Sparse Models and Technology (iTWIST)*, Namur, Belgique, Août 2014.
- [C1] A. Drémeau, L. Daudet – **Structured and soft ! Boltzmann machine and mean-field approximation for structured sparse representations** – *Workshop on Signal Processing with Adaptive Sparse Structured Representations (SPARS)*, Edinburgh, UK, Juin 2011.

Rapports techniques

- [T3] C. Herzet, A. Drémeau – **Joint screening tests for LASSO** – arXiv :1710.09809.
- [T2] A. Drémeau, C. Herzet – **DOA estimation in structured phase-noisy environments : technical report** – arXiv :1609.03503.
- [T1] C. Herzet, A. Drémeau – **Bayesian Pursuit Algorithms** – arXiv :1401.7538.

Thèse de doctorat

- [D] A. Drémeau – **Décompositions parcimonieuses : approches bayésiennes et application à la compression d'image** – Thèse de doctorat de l'Université de Rennes 1, Novembre 2010.

Annexe C

Articles sélectionnés

Dans cette annexe, j'ai sélectionné quelques articles illustrant mes contributions dans le domaine des algorithmes variationnels bayésiens :

- A. Drémeau, C. Herzet, L. Daudet - Boltzmann machine and mean-field approximation for structured sparse decompositions - IEEE Transactions On Signal Processing, Vol. 60, No. 7, pp. 3425-3438, Juillet 2012.

- A. Drémeau, F. Le Courtois, J. Bonnel - Reconstruction of dispersion curves in the frequency-wavenumber domain using compressed sensing on a random array - IEEE Journal of Oceanic Engineering, Vol. 42, No. 4, pp. 914 - 922, Mars 2017.

- C. Dorffer, T. Paviet-Salomon, G. Le Chenadec, A. Drémeau - Modal estimation in underwater acoustics by data-driven structures sparse decomposition - European Signal Processing Conference (EUSIPCO), Virtual conference, Septembre 2021.

- A. Drémeau, A. Liutkus, D. Martina, O. Katz, C. Schülke, F. Krzakala, S. Gigan, L. Daudet - Reference-less measurement of the transmission matrix of a highly scattering material using a DMD and phase retrieval techniques - Optics Express, Vol. 23, pp. 11898-11911, 2015.

- A. Drémeau, C. Herzet - DOA estimation in structured phase-noisy environments - IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Nouvelle Orleans, USA, Mars 2017.

Boltzmann Machine and Mean-Field Approximation for Structured Sparse Decompositions

Angélique Drémeau, Cédric Herzet, and Laurent Daudet, *Senior Member, IEEE*

Abstract—Taking advantage of the structures inherent in many sparse decompositions constitutes a promising research axis. In this paper, we address this problem from a Bayesian point of view. We exploit a Boltzmann machine, allowing to take a large variety of structures into account, and focus on the resolution of a marginalized maximum a posteriori problem. To solve this problem, we resort to a mean-field approximation and the “variational Bayes expectation-maximization” algorithm. This approach results in a soft procedure making no hard decision on the support or the values of the sparse representation. We show that this characteristic leads to an improvement of the performance over state-of-the-art algorithms.

Index Terms—Bernoulli–Gaussian model, Boltzmann machine, mean-field approximation, structured sparse representation.

I. INTRODUCTION

SPARSE representations (SR) aim at describing a signal as the combination of a small number of elementary signals, or atoms, chosen from an overcomplete dictionary. These decompositions have proved useful in a variety of domains including audio [1], [2] and image [3], [4] processing and are at the heart of the recent compressive-sensing paradigm [5].

Formally, let $\mathbf{y} \in \mathbb{R}^N$ be an observed signal and $\mathbf{D} \in \mathbb{R}^{N \times M}$ with $M \geq N$, a dictionary, i.e., a matrix whose columns correspond to atoms. Then one standard formulation of the sparse representation problem can be written as

$$\mathbf{z}^* = \arg \min_{\mathbf{z}} \|\mathbf{y} - \mathbf{D}\mathbf{z}\|_2^2 \quad \text{subject to} \quad \|\mathbf{z}\|_0 \leq L, \quad (1)$$

or, in its Lagrangian version

$$\mathbf{z}^* = \arg \min_{\mathbf{z}} \|\mathbf{y} - \mathbf{D}\mathbf{z}\|_2^2 + \lambda \|\mathbf{z}\|_0, \quad (2)$$

where $\|\mathbf{z}\|_0$ denotes the ℓ_0 pseudo-norm which counts the number of non-zero elements in \mathbf{z} and $L, \lambda > 0$ are parameters specifying the tradeoff between sparsity and distortion.

Manuscript received November 29, 2011; revised March 15, 2012; accepted March 15, 2012. Date of publication April 03, 2012; date of current version June 12, 2012. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Lawrence Carin. The work of A. Drémeau was supported by a fellowship from the Fondation Pierre-Gilles De Gennes pour la Recherche, France.

A. Drémeau is with Institut Langevin, ESPCI ParisTech, CNRS UMR 7587, 75005 Paris, France (e-mail: angelique.dremeau@telecom-paristech.fr).

C. Herzet is with INRIA Centre Rennes-Bretagne Atlantique, 35000 Rennes, France (e-mail: cedric.herzet@inria.fr).

L. Daudet is with Institut Langevin, ESPCI ParisTech, CNRS UMR 7587, 75005 Paris, France, on a joint affiliation with Université Paris Diderot-Paris 7 and the Institut Universitaire de France (e-mail: laurent.daudet@espci.fr).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2012.2192436

Finding the exact solution of (1), (2) is an NP-hard problem [6], i.e., it generally requires a combinatorial search over the entire solution space. Therefore, heuristic (but tractable) algorithms have been devised to deal with this problem. These algorithms are based on different strategies that we review in Section I-A.

More recently, the SR problem has been enhanced by the introduction of structural constraints on the support of the sparse representation: the non-zero components of \mathbf{z} can no longer be chosen independently from each other but must obey some (deterministic or probabilistic) inter-rules. This problem is often referred to as “structured” sparse representation. This new paradigm has been found to be relevant in many application domains and has recently sparked a surge of interest in algorithms dealing with this problem (see Section I-B).

In this paper, we propose a novel algorithm addressing the SR problem in the structured setup and consider the standard, non-structured setup as a particular case. The proposed algorithm is cast within a Bayesian inference framework and based on the use of a particular variational approximation as a surrogate to an optimal maximum a posteriori (MAP) decision. In order to properly place our work in the rich literature pertaining to SR algorithms, we briefly review hereafter some of the algorithms coping with the standard and the structured SR problems.

A. Standard Sparse Representation Algorithms

The origin of the algorithms addressing the standard sparse representation problem (1), (2) traces back to the fifties, e.g., in the field of statistical regression [7] and operational research [8], [9]. The algorithms available today in the literature can roughly be divided into four main families:

- 1) *The algorithms based on problem relaxation*: these procedures replace the ℓ_0 -norm by an ℓ_p -norm (with $0 < p \leq 1$). This approximation leads to a relaxed problem which can be solved efficiently by standard optimization procedures. Well-known instances of algorithms based on such an approach are the Basis Pursuit (BP) [10], Least Absolute Shrinkage and Selection Operator (LASSO) [11] or Focal Underdetermined System Solver (FocUSS) [12] algorithms.
- 2) *The iterative thresholding algorithms*: these procedures build up the sparse vector \mathbf{z} by making a succession of thresholding operations. The first relevant work in this family was realized by Kingsbury and Reeves [13] who derive an iterative thresholding method with the aim at solving problem (2). However, their contribution is done without a clear connection to the objective function (2). We find a more explicit version of their results in [14]

where Blumensath and Davies introduce the Iterative Hard Thresholding (IHT) algorithm. Daubechies *et al.* propose in [15] a similar procedure while replacing the ℓ_0 -norm by the ℓ_1 -norm. The resulting algorithm relies then on a soft thresholding operation.

- 3) *The pursuit algorithms*: these methods build up the sparse vector \mathbf{z} by making a succession of greedy decisions. There exist many pursuit algorithms in the current literature. Among the most popular, we can cite Matching Pursuit (MP) [16], Orthogonal Matching Pursuit (OMP) [17] or Orthogonal Least Square (OLS) [18]. The latter algorithms do not allow for the selection of more than one atom per iteration. This limitation is avoided by more recent procedures like Stagewise OMP (StOMP) [19], Subspace Pursuit (SP) [20] or Compressive Sampling Matching Pursuit (CoSaMP) [21].
- 4) *The Bayesian algorithms*: these procedures express the SR problem as the solution of a Bayesian inference problem and apply statistical tools to solve it. They mainly distinguish by the prior model, the considered estimation problem and the type of statistical tools they apply to solve it. Regarding the choice of the prior, a popular approach consists in modeling \mathbf{z} as a continuous random variable whose distribution has a sharp peak at zero and heavy tails (e.g., Cauchy [22], Laplace [23], [24], *t*-Student [25], Jeffrey's [26] distributions). Another approach, recently gaining in popularity, is based on a prior made up of the combination of Bernoulli and Gaussian distributions [27]–[36]. Different variants of Bernoulli–Gaussian (BG) models exist. A first approach, as considered in [27], [30], [31], [34], consists in assuming that the elements of \mathbf{z} are independently drawn from Gaussian distributions whose variances are controlled by Bernoulli variables: a small variance enforces elements to be close to zero whereas a large one defines a non-informative prior on non-zero coefficients. Another model on \mathbf{z} based on BG variables is as follows: the elements of the sparse vector are defined as the multiplication of Gaussian and Bernoulli variables. This model has been exploited in the contributions [28], [29], [32], [33] and will be considered in the present paper. These two distinct hierarchical BG models share a similar marginal expression of the form:

$$p(\mathbf{z}) = \prod_{i=1}^M (p_i \mathcal{N}(0, \sigma_0^2) + (1 - p_i) \mathcal{N}(0, \sigma_1^2)) \quad (3)$$

where the p_i 's are the parameters of the Bernoulli variables. While σ_0^2 can be tuned to any positive real value in the first BG model presented above, it is set to 0 in the second one. This marginal formulation is directly used in many contributions as in [35] and [36].

B. Structured Sparse Representation Algorithms

The algorithms dedicated to “standard” SR problems (1), (2) do not assume any dependency between the non-zero elements of the sparse vector, i.e., they select the atoms of the sparse decomposition without any consideration of possible

links between them. Yet, recent contributions have shown the existence of structures in many natural signals (depending on the dictionary and the class of signals) and emphasize the relevance of exploiting them in the process of sparse decomposition. Hence, many contributions have recently focused on the design of “structured” sparse representation algorithms, namely algorithms taking the dependencies between the elements of SR support into account.

The algorithms available in the literature essentially rely on the same type of approximation as their standard counterpart (see Section I-A) and could be classified accordingly. We found however more enlightening to present the state-of-the-art contributions according to the type of structure they exploit. We divide them into four families:

- 1) *Group sparsity*: in group-sparse signals, coefficients are either all non-zero or all zero within pre-specified groups of atoms. This type of structure is also referred to as *block sparsity* in some contributions [37], [38]. In practice, group sparsity can be enforced by the use of particular “mixed” norms combining ℓ_1 - and ℓ_2 -norms. Following this approach, Yuan and Lin propose in [39] a LASSO-based algorithm called Group-LASSO, while in [37], Eldar and Mishali derive a modified SOCP (Second Order Cone Program) algorithm and in [38], Eldar *et al.* introduce the Block-OMP, group-structured extension of OMP. Parallel to these contributions, other approaches have been proposed. Let us mention [40] and [41] based on clusters, [42] where coding costs are considered, or [43] relying on the definition of Boolean variables and the use of an approximate message passing algorithm [44]. Finally, as an extension of group sparsity, Sprechmann *et al.* consider in [45] intra-group sparsity by means of an additional penalty term.
- 2) *Molecular sparsity*: molecular sparsity describes more complex structures, in the particular case where the atoms of the dictionary have a double indexing (e.g., time-frequency atoms). It can be seen as the combination of two group-sparsity constraints: one on each component of the double index. This type of structure is also referred to as *elitist sparsity* by certain authors [46]. In order to exploit molecular sparsity, Kowalski and Torr sani study in [46] the general use of mixed norms in structured sparsity problems. They thus motivate the Group-LASSO algorithm introduced in [39] and propose an extension of it, the Elitist-LASSO. Molecular sparsity has also been considered by Daudet in [47] for audio signals: the paper introduces the Molecular-MP algorithm which uses a local tonality index.
- 3) *Chain and tree-structured sparsity*: such structures arise in many applications. For example, chain structure appears in any sequential process whereas tree-structured sparsity is at the heart of wavelet decompositions, widely used in image processing. *De facto*, we find in the literature several contributions dealing with these particular types of constrained sparsity. Tree-structured sparsity is addressed in [48] where the authors define a particular penalty term replacing the commonly used ℓ_0 - or ℓ_1 -norms, and [49], [50] which define a probabilistic framework based on

Bernoulli variables with scale-depending parameters. These two latter contributions focus on the sampling of the posterior distribution of \mathbf{z} and resort either to Monte Carlo Markov chain (MCMC) methods or to mean-field approximations. Chain-structured sparsity can be enforced using a Markov-chain process. This is for example the model adopted by Févotte *et al.* in [2], combined then with a MCMC inference scheme, and by Schniter in [51], together with an approximate message passing algorithm [44].

- 4) *Generic structured sparsity*: some approaches do not focus on a specific type of structure but propose general models accounting for a wide set of structures. Most of these approaches are probabilistic. In particular, [52]–[54] have recently emphasized the relevance of the Boltzmann machine as a general model for structured sparse representations. Well-known in Neural Networks, this model allows indeed to consider dependencies between distant atoms and thus constitutes an adaptive framework for the design of structured SR algorithms. In this paper, we will consider this particular model to derive a novel structured SR algorithm.

Finally, let us mention the deterministic approach in [55] which introduces the model-based CoSaMP, relying on the definition of a “model” peculiar to a structure. As practical examples, the authors apply their algorithm to group and tree-structured sparsity.

C. Contributions of This Paper

In this paper, we focus on the design of an effective structured SR algorithm within a Bayesian framework. Motivated by a previous result [32], a Boltzmann machine is introduced to describe general sparse structures. In this context, we reformulate the structured sparse representation problem as a particular marginalized maximum *a posteriori* (MAP) problem on the support of the sparse vector. We then apply a particular variational mean-field approximation to deal with the intractability of the original problem; this results in the so-called “SSoBaP” algorithm. We emphasize that SSoBaP shares some structural similarities with MP but enjoys additional desirable features: i) it can exploit a number of different structures on the support and ii) its iterative process is based on the exchange of *soft* decisions (by opposition to *hard* decisions for MP) on the support. We confirm through simulation results that SSoBaP leads to an improvement of the reconstruction performance (according to several figures of merits) over several SR algorithms of the state-of-the-art.

D. Organization of This Paper

The paper is organized as follows. Section II describes the probabilistic model used to derive our algorithm. In particular, we suppose that the SR support is distributed according to a Boltzmann machine and show that this model allows to describe many well-known probabilistic model as particular cases. In this framework, Section III presents different Bayesian estimators which can be considered within the SR problematic. We focus in particular on a marginalized maximum *a posteriori* (MAP) problem on the SR support.

Section IV is dedicated to the resolution of this MAP problem. We propose in this paper to resort to a mean-field approximation and the “variational Bayes Expectation-Maximization” algorithm. The first subsection of Section IV recalls the basics of this variational approach. The rest of the section is dedicated to the description of the proposed algorithm.

The performance of the proposed algorithm is evaluated in Section V by various experiments involving different evaluation criteria on synthetic data. We show that, as long as our simulation setups are concerned, the proposed algorithm is very competitive with state-of-the-art procedures.

II. PROBABILISTIC MODEL

Let $\mathbf{x} \in \mathbb{R}^M$ be a vector defining the amplitudes of the sparse representation and $\mathbf{s} \in \{0, 1\}^M$ be a vector defining the SR support, i.e., the subset of columns of \mathbf{D} used to generate \mathbf{y} . Without loss of generality, we will adopt the following convention: if $s_i = 1$ (resp. $s_i = 0$), the i th column of \mathbf{D} is (resp. is not) used to form \mathbf{y} . Denoting by \mathbf{d}_i the i th column of \mathbf{D} , we then consider the following observation model¹:

$$\mathbf{y} = \sum_{i=1}^M s_i x_i \mathbf{d}_i + \mathbf{n}, \quad (4)$$

where \mathbf{n} is a zero-mean white Gaussian noise with variance σ_n^2 . Therefore,

$$p(\mathbf{y}|\mathbf{x}, \mathbf{s}) = \mathcal{N}(\mathbf{D}_s \mathbf{x}_s, \sigma_n^2 \mathbf{I}_N), \quad (5)$$

where \mathbf{I}_N is the $N \times N$ -identity matrix and \mathbf{D}_s (resp. \mathbf{x}_s) is a matrix (resp. vector) made up of the \mathbf{d}_i 's (resp. x_i 's) such that $s_i = 1$. We suppose that \mathbf{x} obeys the following probabilistic model:

$$p(\mathbf{x}) = \prod_{i=1}^M p(x_i) \quad \text{where} \quad p(x_i) = \mathcal{N}(0, \sigma_{x_i}^2). \quad (6)$$

Within model (5), (6), the observation \mathbf{y} is thus seen as the noisy combination of atoms specified by \mathbf{s} . The weights of the combination are realizations of Gaussian distributions whose variances are independent on the support \mathbf{s} .

Clearly both the number of atoms building up \mathbf{y} as well as their interdependencies are a function of the prior defined on \mathbf{s} . A standard choice for modelling *unstructured* sparsity is based on a product of Bernoulli distributions, i.e.,

$$p(\mathbf{s}) = \prod_{i=1}^M p(s_i) \quad \text{where} \quad p(s_i) = \text{Ber}(p_i), \quad (7)$$

and $p_i \ll 1$. This model is indeed well-suited to modelling situations where \mathbf{y} stems from a sparse process: if $p_i \ll 1 \forall i$, only a small number of s_i 's will *typically*² be non-zero, i.e., the observation vector \mathbf{y} will be generated with high probability from a small subset of the columns of \mathbf{D} . In particular, if $p_i = p$

¹The sparse representation \mathbf{z} , as used in Section I, is then defined as the Hadamard product of \mathbf{x} and \mathbf{s} , i.e., $z_i = s_i x_i, \forall i \in \llbracket 1, M \rrbracket$.

²In an information-theoretic sense, i.e., according to model (5)–(7), a realization of \mathbf{s} with a few non-zero components will be observed with probability almost 1.

$\forall i$, typical realizations of \mathbf{y} will involve a combination of pM columns of \mathbf{D} [see (8)–(9) at the bottom of the page].

Note that (7) does not impose any interaction between the atoms building up the observed vector \mathbf{y} : each s_i is the realization of an independent random variable. Taking atom interdependencies into account therefore requires more involved probabilistic models. The so-called *Boltzmann machine* offers a nice option for this purpose [56]. Formally, it can be expressed as

$$p(\mathbf{s}) \propto \exp(\mathbf{b}^T \mathbf{s} + \mathbf{s}^T \mathbf{W} \mathbf{s}), \quad (10)$$

where \mathbf{W} is a symmetric matrix with zeros on the diagonal and \propto denotes equality up to a normalization factor. Parameter \mathbf{b} defines the biases peculiar to each element of \mathbf{s} while $\mathbf{W} \in \mathbb{R}^{M \times M}$ characterizes the interactions between them: w_{ij} weights the dependency between atoms s_i and s_j .

The Boltzmann machine encompasses many well-known probabilistic models as particular cases. For example, the Bernoulli model (7) corresponds to $\mathbf{W} = \mathbf{0}_{M \times M}$ (expressing the atoms' independence):

$$p(\mathbf{s}) \propto \exp(\mathbf{b}^T \mathbf{s}) = \prod_i \exp(b_i s_i), \quad (11)$$

which is equivalent to a Bernoulli model (7) with

$$p_i = \frac{1}{1 + \exp(-b_i)}. \quad (12)$$

Another example is the Markov chain. For instance, let us consider the following first-order Markov chain:

$$p(\mathbf{s}) = p(s_1) \prod_{i=1}^M p(s_{i+1} | s_i), \quad (13)$$

with $\forall i \in \llbracket 1, M-1 \rrbracket$,

$$p(s_{i+1} = s | s_i = 1) \triangleq \begin{cases} 1 - p_{i+1}^1 & \text{if } s = 1, \\ p_{i+1}^1 & \text{if } s = 0, \end{cases} \quad (14)$$

$$p(s_{i+1} = s | s_i = 0) \triangleq \begin{cases} p_{i+1}^0 & \text{if } s = 1, \\ 1 - p_{i+1}^0 & \text{if } s = 0, \end{cases} \quad (15)$$

$$p(s_1 = s) \triangleq \begin{cases} p_1 & \text{if } s = 1, \\ 1 - p_1 & \text{if } s = 0. \end{cases} \quad (16)$$

This Markov chain corresponds to a Boltzmann machine with parameters \mathbf{b} and \mathbf{W} defined in (8), (9). In particular, only two subdiagonals in \mathbf{W} are non-zero.

In the rest of this paper, we will derive the main equations of our algorithm for the general model (10). We will then particularize them to model (7), which leads to an algorithm for standard (unstructured) sparse representation.

III. SPARSE REPRESENTATIONS WITHIN A BAYESIAN FRAMEWORK

The probabilistic framework defined in Section II allows us to tackle the SR problem from a Bayesian perspective. As long as (5), (6) is the true generative model for the observations \mathbf{y} , optimal estimators can be derived under different Bayesian criteria (mean square error, mean absolute error, etc.). We focus hereafter on the computation of a solution under a MAP criterion, which corresponds to the optimal Bayesian estimator for a Bayesian cost based on a “notch” loss function [57].

A first possible approach consists in solving the *joint* MAP problem:

$$(\hat{\mathbf{x}}, \hat{\mathbf{s}}) = \arg \max_{\mathbf{x}, \mathbf{s}} \log p(\mathbf{x}, \mathbf{s} | \mathbf{y}). \quad (17)$$

Interestingly, we emphasize in [32] that the joint MAP problem (17) shares the same set of solutions as the standard SR problem (2) within BG model (6), (7). This connection builds a bridge between standard and Bayesian SR procedures and motivates the use of model (6), (7) (and its structured generalization (6)–(10)) in other estimation problems. In particular, we focus hereafter on MAP problems oriented to the recovery of the SR support.

Assuming (5), (6) is the true generative model of \mathbf{y} , the decision minimizing the probability of wrong decision on the *whole* SR support is given by

$$\hat{\mathbf{s}} = \arg \max_{\mathbf{s} \in \{0,1\}^M} \log p(\mathbf{s} | \mathbf{y}), \quad (18)$$

where $p(\mathbf{s} | \mathbf{y}) = \int_{\mathbf{x}} p(\mathbf{x}, \mathbf{s} | \mathbf{y}) d\mathbf{x}$. Problem (18) is unfortunately intractable since it typically requires the evaluation of the cost function, $\log p(\mathbf{s} | \mathbf{y})$, for all possible 2^M sequences in $\{0, 1\}^M$. A heuristic greedy procedure looking for the solution of (18) has recently been proposed in [54].

In this paper, we address the SR representation problem from a different perspective. The decision on each element of the support is made from a marginalized MAP estimation problem:

$$\hat{s}_i = \arg \max_{s_i \in \{0,1\}} \log p(s_i | \mathbf{y}). \quad (19)$$

$$\mathbf{b} = \left[\log \frac{p_1}{1-p_1} + \log \frac{p_2^1}{1-p_2^0} \dots \log \frac{p_i^0}{1-p_i^0} + \log \frac{p_{i+1}^1}{1-p_{i+1}^1} \dots \log \frac{p_M^0}{1-p_M^0} \right]^T, \quad (8)$$

$$\mathbf{W} = \begin{pmatrix} 0 & \frac{1}{2} \log \frac{(1-p_2^1)(1-p_2^0)}{p_2^1 p_2^0} & 0 & 0 \\ \frac{1}{2} \log \frac{(1-p_2^1)(1-p_2^0)}{p_2^1 p_2^0} & 0 & \frac{1}{2} \log \frac{(1-p_3^1)(1-p_3^0)}{p_3^1 p_3^0} & \vdots \\ 0 & \frac{1}{2} \log \frac{(1-p_3^1)(1-p_3^0)}{p_3^1 p_3^0} & \ddots & \vdots \\ \vdots & \ddots & 0 & \vdots \\ 0 & \dots & \dots & 0 \end{pmatrix}. \quad (9)$$

The solution of (19) minimizes the probability of making a wrong decision on *each* s_i (rather than on the whole sequence as in (18)).

At first sight, problem (19) may appear easy to solve since the search space only contains two elements i.e., $s_i \in \{0, 1\}$. However, the evaluation of $p(s_i|\mathbf{y})$ turns out to be intractable since it requires a costly marginalization of the joint probability $p(\mathbf{s}|\mathbf{y})$ over the s_j 's, $j \neq i$. Nevertheless, many tools exist in the literature to circumvent this issue. In particular, the family of variational approximations allows for the computation of tractable surrogates of $p(s_i|\mathbf{y})$, say $q(s_i)$; see [58]. In this paper we will resort to a *mean-field* variational approximation to compute a surrogate $q(s_i)$ of $p(s_i|\mathbf{y})$ (see Section IV). In particular, in this paper we will resort to a mean-field variational approximation to compute a tractable surrogate of $p(s_i|\mathbf{y})$, say $q(s_i)$. Problem (19) will then be approximated by

$$\hat{s}_i = \arg \max_{s_i \in \{0,1\}} \log q(s_i), \quad (20)$$

which is straightforward to solve.

Finally, given the estimated support $\hat{\mathbf{s}}$, we can reconstruct the coefficients of a sparse representation say $\hat{\mathbf{x}}_{\hat{\mathbf{s}}}$, as its MAP estimate

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x}} \log p(\mathbf{x}|\hat{\mathbf{s}}, \mathbf{y}). \quad (21)$$

The solution of (21) is expressed as

$$\begin{aligned} \hat{\mathbf{x}}_{\hat{\mathbf{s}}} &= (\mathbf{D}_{\hat{\mathbf{s}}}^T \mathbf{D}_{\hat{\mathbf{s}}} + \Delta)^{-1} \mathbf{D}_{\hat{\mathbf{s}}}^T \mathbf{y}, \\ \text{and } \hat{x}_i &= 0 \quad \text{if } s_i = 0, \end{aligned} \quad (22)$$

where Δ is a diagonal matrix whose i th element is $\frac{\sigma_x^2}{\sigma_{x_i}^2}$. When $\sigma_{x_i}^2 \rightarrow +\infty \forall i$, (22) reduces to the least-square estimate

$$\begin{aligned} \hat{\mathbf{x}}_{\hat{\mathbf{s}}} &= \mathbf{D}_{\hat{\mathbf{s}}}^+ \mathbf{y}, \\ \text{and } \hat{x}_i &= 0 \quad \text{if } s_i = 0, \end{aligned} \quad (23)$$

where $\mathbf{D}_{\hat{\mathbf{s}}}^+$ is the Moore–Penrose pseudo-inverse of the matrix made up of the \mathbf{d}_i 's such that $\hat{s}_i = 1$.

IV. STRUCTURED SOFT BAYESIAN PURSUIT ALGORITHM

In this section, we detail our methodology to compute the approximation $q(s_i)$ of the posterior probability $p(s_i|\mathbf{y})$. Our approach is based on a well-known variational approximation, namely the mean-field (MF) approximation, and its practical implementation via the so-called VB-EM algorithm. This methodology results in an iterative algorithm whose updates are very similar to those of the recently-proposed Bayesian Matching Pursuit algorithm (BMP) [32]. However, unlike the latter, the proposed procedure updates probabilities rather than estimates of the SR support. Moreover, BMP as introduced in [32] does not deal with structured sparsity. In the sequel, we will thus refer to the proposed procedure as the ‘‘Structured Soft Bayesian Pursuit algorithm’’ (SSoBaP).

The rest of this section is organized as follows. We first briefly recall the general theory pertaining to mean-field approximations. Then, in Section IV-B we derive the main equations defining SSoBaP. The Section IV-C is dedicated to the Soft Bayesian Pursuit algorithm (SoBaP), particular case of SSoBaP resulting from the choice $\mathbf{W} = \mathbf{0}_{M \times M}$ in the Boltzmann machine (10). In the next subsection, we emphasize the

advantage of making soft decisions by comparing the update equations of BMP and SSoBaP. We address the problem of parameter estimation in a ‘‘variational Bayes Expectation-Maximization’’ framework in Section IV-E and finally, emphasize the differences and connections of SSoBaP (and SoBaP) with existing algorithms in the last subsection.

A. Mean-Field Approximation: Basics

The mean-field approximation [59] refers to a family of approximations of posterior probabilities by distributions having a ‘‘tractable’’ factorization. Formally, let $\boldsymbol{\theta}$ denote a vector of random variables and $p(\boldsymbol{\theta}|\mathbf{y})$ its a posteriori probability. Let moreover $(\boldsymbol{\theta}_i)_{i=1}^I$ denotes a *partition* of the elements of $\boldsymbol{\theta}$ i.e.,

$$\boldsymbol{\theta} = [\boldsymbol{\theta}_1^T \dots \boldsymbol{\theta}_I^T]^T. \quad (24)$$

Then, the mean-field approximation of $p(\boldsymbol{\theta}|\mathbf{y})$ relative to partition (24) is the surrogate distribution $q^*(\boldsymbol{\theta})$ satisfying

$$q^*(\boldsymbol{\theta}) = \arg \min_{q(\boldsymbol{\theta})} \left\{ \int q(\boldsymbol{\theta}) \log \left(\frac{q(\boldsymbol{\theta})}{p(\boldsymbol{\theta}|\mathbf{y})} \right) d\boldsymbol{\theta} \right\}, \quad (25)$$

subject to

$$q(\boldsymbol{\theta}) = \prod_{i=1}^I q(\boldsymbol{\theta}_i), \quad \int_{\boldsymbol{\theta}_i} q(\boldsymbol{\theta}_i) d\boldsymbol{\theta}_i = 1 \quad \forall i \in \llbracket 1, I \rrbracket. \quad (26)$$

The mean-field approximation $q^*(\boldsymbol{\theta})$ is therefore the distribution minimizing the Kullback–Leibler divergence with the actual posterior $p(\boldsymbol{\theta}|\mathbf{y})$ while factorizing as a product of probabilities (26). There potentially³ are as many possible mean-field approximations as partitions of $\boldsymbol{\theta}$. In practice, the choice of a particular approximation results from a tradeoff between complexity and accuracy.

A solution to problem (25), (26) can be looked for by successively minimizing the Kullback–Leibler divergence with respect to one single factor, say $q(\boldsymbol{\theta}_i)$. This gives rise to the following update equations

$$\begin{aligned} q^{(n+1)}(\boldsymbol{\theta}_1) &\propto \exp \left\{ \langle \log p(\boldsymbol{\theta}, \mathbf{y}) \rangle_{\prod_{j \neq 1} q^{(n)}(\boldsymbol{\theta}_j)} \right\}, \\ &\vdots \\ q^{(n+1)}(\boldsymbol{\theta}_i) &\propto \exp \left\{ \langle \log p(\boldsymbol{\theta}, \mathbf{y}) \rangle_{\prod_{j > i} q^{(n)}(\boldsymbol{\theta}_j) \prod_{j < i} q^{(n+1)}(\boldsymbol{\theta}_j)} \right\}, \\ &\vdots \\ q^{(n+1)}(\boldsymbol{\theta}_I) &\propto \exp \left\{ \langle \log p(\boldsymbol{\theta}, \mathbf{y}) \rangle_{\prod_{j \neq I} q^{(n+1)}(\boldsymbol{\theta}_j)} \right\}, \end{aligned} \quad (27)$$

where

$$\langle \log p(\boldsymbol{\theta}, \mathbf{y}) \rangle_{q(\boldsymbol{\theta}_i)} \triangleq \int_{\boldsymbol{\theta}_i} q(\boldsymbol{\theta}_i) \log p(\boldsymbol{\theta}, \mathbf{y}) d\boldsymbol{\theta}_i. \quad (28)$$

Note that we suppose in (27) that the $q(\boldsymbol{\theta}_i)$'s are updated at each iteration one after the other, in an increasing order of their

³Two different partitions can indeed lead incidentally to the same solution for (25), (26).

indices. However the extension to other update schedulings is straightforward.

The procedure described in (27) is usually referred to as *variational Bayes Expectation-Maximization (VB-EM) algorithm* in the literature [60]–[62]. VB-EM is ensured to converge to a saddle point or a (local or global) maximum of problem (25), (26) under mild conditions.

The appellation “VB-EM” comes from the close connection of the above procedure with the well-known EM algorithm [63]. The relation between the two algorithms can be seen by imposing an additional constraint on some $q(\boldsymbol{\theta}_i)$'s, namely

$$q(\boldsymbol{\theta}_i) = \delta(\boldsymbol{\theta}_i - \hat{\boldsymbol{\theta}}_i), \quad (29)$$

where $\delta(\cdot)$ denotes the Dirac delta function. Minimizing the Kullback–Leibler divergence with respect to $q(\boldsymbol{\theta}_i)$ while taking (29) into account then reduces to optimizing the value of $\hat{\boldsymbol{\theta}}_i$. Thus, for the $\boldsymbol{\theta}_i$'s subject to (29), the update (27) can be rewritten as

$$\hat{\boldsymbol{\theta}}_i^{(n+1)} = \arg \max_{\boldsymbol{\theta}_i} \left\{ \langle \log p(\boldsymbol{\theta}, \mathbf{y}) \rangle_{\prod_{j>i} q^{(n)}(\boldsymbol{\theta}_j)} \prod_{j<i} q^{(n+1)}(\boldsymbol{\theta}_j) \right\}. \quad (30)$$

Now, let $\boldsymbol{\theta}_{\neq i}$ denote the vector made of the $\boldsymbol{\theta}_j$'s, $j \neq i$. If only one element in the partition $(\boldsymbol{\theta}_i)_{i=1}^I$, say $\boldsymbol{\theta}_j$, is *not* subject to (29), it can be shown [64] that the update (27)–(30) define an EM algorithm aiming at solving

$$\hat{\boldsymbol{\theta}}_{\neq j} = \arg \max_{\boldsymbol{\theta}_{\neq j}} \log p(\boldsymbol{\theta}_{\neq j} | \mathbf{y}), \quad (31)$$

where $\boldsymbol{\theta}_j$ is considered as a hidden variable. The E-step then corresponds to the estimation of $q(\boldsymbol{\theta}_j)$ (27), namely $p(\boldsymbol{\theta}_j | \mathbf{y}, \boldsymbol{\theta}_{\neq j})$ in this particular case, while the M-step computes (30) $\forall i \neq j$, i.e., maximizes expectation $E_{\boldsymbol{\theta}_j} [\log p(\boldsymbol{\theta}, \mathbf{y})]$ with respect to parameters $\hat{\boldsymbol{\theta}}_{\neq j}$.

The general case where several $\boldsymbol{\theta}_i$'s are not subject to (29) (as opposed to the case presented above where all $\boldsymbol{\theta}_i$ but one where subject to (29)) does not correspond to an EM algorithm anymore as the E-step does not reduce to the estimation of *one* posterior probability but *approximates* a joint probability by means of an MF approximation.

To conclude this section, let us point out that mean-field approximations offer a nice framework to approximate the marginals $p(\boldsymbol{\theta}_i | \mathbf{y})$'s, where $\boldsymbol{\theta}_i$ is an element of the mean field partition (24) (note that we use here the word “marginal” in a large sense since $\boldsymbol{\theta}_i$ possibly contains more than one variable). Indeed, assume one wants to compute

$$p(\boldsymbol{\theta}_i | \mathbf{y}) = \int_{\boldsymbol{\theta}_{\neq i}} p(\boldsymbol{\theta} | \mathbf{y}) d\boldsymbol{\theta}_{\neq i}. \quad (32)$$

Then, using the decomposition property of the mean-field approximation (26), we come up with

$$p(\boldsymbol{\theta}_i | \mathbf{y}) \simeq \int_{\boldsymbol{\theta}_{\neq i}} q(\boldsymbol{\theta}) d\boldsymbol{\theta}_{\neq i} \quad (33)$$

$$\simeq q(\boldsymbol{\theta}_i) \int_{\boldsymbol{\theta}_{\neq i}} \prod_{j \neq i} q(\boldsymbol{\theta}_j) d\boldsymbol{\theta}_{\neq i} = q(\boldsymbol{\theta}_i). \quad (34)$$

Factors $q(\boldsymbol{\theta}_i)$'s are therefore approximations of marginals. We will exploit this observation in the next section to derive a tractable approximation of $p(s_i | \mathbf{y})$.

B. SSoBaP

In this paper, we consider the particular case where the MF approximation of $p(\mathbf{x}, \mathbf{s} | \mathbf{y})$, say $q(\mathbf{x}, \mathbf{s})$, is constrained to have the following structure:

$$q(\mathbf{x}, \mathbf{s}) = \prod_i q(x_i, s_i). \quad (35)$$

This is equivalent to setting $I = M$, $\boldsymbol{\theta} = [x_1 s_1 \dots x_M s_M]^T$ and $\boldsymbol{\theta}_i = [x_i s_i]^T \forall i$ in the general framework described in Section IV-A. Note that, the $\boldsymbol{\theta}_i$'s do not correspond to single elements of $\boldsymbol{\theta}$ but form a partition of $\boldsymbol{\theta}$.

Particularized to model (5), (6)–(10), the corresponding VB-EM update (27) are written as⁴:

$$q(x_i, s_i) = q(x_i | s_i) q(s_i), \quad (36)$$

where

$$q(x_i | s_i) = \mathcal{N}(m(s_i), \Sigma(s_i)), \quad (37)$$

$$q(s_i) \propto \sqrt{\Sigma(s_i)} \exp\left(\frac{1}{2} \frac{m(s_i)^2}{\Sigma(s_i)}\right) \times \exp\left(s_i \left(b_i + 2 \sum_{j \neq i} q(s_j = 1) w_{ij}\right)\right), \quad (38)$$

and

$$\Sigma(s_i) = \frac{\sigma_{x_i}^2 \sigma_n^2}{\sigma_n^2 + s_i \sigma_{x_i}^2 \mathbf{d}_i^T \mathbf{d}_i}, \quad (39)$$

$$m(s_i) = s_i \frac{\sigma_{x_i}^2}{\sigma_n^2 + s_i \sigma_{x_i}^2 \mathbf{d}_i^T \mathbf{d}_i} \langle \mathbf{r}_i \rangle^T \mathbf{d}_i, \quad (40)$$

$$\langle \mathbf{r}_i \rangle = \mathbf{y} - \sum_{j \neq i} q(s_j = 1) m(s_j = 1) \mathbf{d}_j. \quad (41)$$

After convergence of the procedure defined in (36)–(41), probabilities $q(x_i, s_i)$ correspond to a mean-field approximation of $p(x_i, s_i | \mathbf{y})$ (see (34)). Coming back to problem (19), an approximation of $p(s_i | \mathbf{y})$ thus simply follows from the relations:

$$p(s_i | \mathbf{y}) = \int p(x_i, s_i | \mathbf{y}) dx_i, \quad (42)$$

$$\simeq \int q(x_i, s_i) dx_i = q(s_i). \quad (43)$$

This approximation can be used in problem (20) to make an approximated MAP decision on s_i . Note that (20) is easy to solve by simple thresholding operation, i.e., $\hat{s}_i = 1$ if $q(s_i = 1) > T$ and $\hat{s}_i = 0$ otherwise, with $T = 0.5$.

The most expensive operation is the update (41) which scales as $\mathcal{O}(NM)$. So, the complexity of one update step is equal to Matching Pursuit (MP). However, in MP *one* unique couple (x_i, s_i) is involved at each iteration while in the proposed algorithm *all* indices are updated one after the other. To the extent of our experiments (see Section V), we could observe that the

⁴When clear from the context, we will drop the iteration indices for notational simplicity.

proposed algorithm converges in a reasonable number of iterations, keeping it at a competitive position beside state-of-the-art algorithms.

C. A Particular Case: SoBaP

As emphasized in Section II, the Boltzmann machine can be seen as a general framework including a large set of probabilistic models. Among them, the Bernoulli model (7) is of particular interest, as a possible approach to model unstructured sparsity (see, e.g., [32], [33]).

From a mathematical point of view, the Bernoulli model (7) corresponds to the simple case $\mathbf{W} = \mathbf{0}_{M \times M}$ in the Boltzmann machine (10). In this case, procedure (36)–(41) remains unchanged, except for (38) which becomes

$$q(s_i) \propto \sqrt{\Sigma(s_i)} \exp\left(\frac{1}{2} \frac{m(s_i)^2}{\Sigma(s_i)}\right) p(s_i), \quad (44)$$

with $p(s_i) = \text{Ber}(p_i)$, $\forall i \in \llbracket 1, M \rrbracket$.

As the BG model (6), (7) is largely used to address the unstructured SR problem, it is useful to distinguish the procedure using (44) from the SSoBaP process. To this end, we will refer to this particular case as “Soft Bayesian Pursuit algorithm” (SoBaP) in the sequel. Note that SoBaP was introduced from a BG perspective in our conference paper [65].

D. Soft Versus Hard Decision

Contrarily to many deterministic (e.g., [55] for structured sparsity [10], [16], [17], [19] for unstructured sparsity) and probabilistic (e.g., [52]–[54] for structured sparsity [32], [33], [66] for unstructured sparsity) algorithms in the literature, the procedure defined in (36)–(41) does not make any hard decision on the SR support or the values of the SR coefficients at each iteration, but evaluates probabilities. It thus allows, to some extent, to take into account the uncertainties we have on the model and to refine this model at each iteration before making the final decision. In particular, it is worth comparing the proposed procedure to the Bayesian Matching Pursuit (BMP) introduced in [32] for unstructured sparsity.

BMP is an iterative procedure looking sequentially for a solution of (17). It proceeds like its standard homologue MP by modifying one unique couple (x_i, s_i) at each iteration, namely the one leading to the highest increase of $\log p(\mathbf{x}, \mathbf{s}|\mathbf{y})$. It can then be shown (see [32]) that the (locally) optimal update of the selected coefficient x_i is given by

$$\hat{x}_i^{(n)} = \hat{s}_i^{(n)} \frac{\sigma_{x_i}^2}{\sigma_n^2 + \sigma_{x_i}^2 \mathbf{d}_i^T \mathbf{d}_i} \mathbf{r}_i^{(n)T} \mathbf{d}_i, \quad (45)$$

$$\text{where } \mathbf{r}_i^{(n)} = \mathbf{y} - \sum_{j \neq i} \hat{s}_j^{(n-1)} \hat{x}_j^{(n-1)} \mathbf{d}_j, \quad (46)$$

and n is the iteration number. We omit here deliberately the support update, addressed in BMP from an “unstructured” point of view.

BMP and SSoBaP share some similarities. In particular, the mean of distribution $q(x_i|s_i)$ computed by the proposed algorithm (40) has the same form as the coefficient update performed by BMP (45). They rely however on different variables, namely the residual \mathbf{r}_i , (46), and its mean $\langle \mathbf{r}_i \rangle$, (41). This fundamental

difference between both algorithms leads to well distinct approaches. In BMP, a hard decision is made on the SR support at each iteration: the atoms of the dictionary are either used or not (each $\hat{x}_j^{(n-1)}$ is multiplied by $\hat{s}_j^{(n-1)}$ which is equal to 0 or 1). On the contrary, in the proposed algorithm, the contributions of the atoms are simply weighted by $q(s_j = 1)$, i.e., the probability distributions of the s_j 's. In a similar way, the coefficients $\hat{x}_j^{(n-1)}$'s used in (46) are replaced by their means $m(s_j = 1)$ in (41), taking into account the uncertainties we have on the values of the x_j 's.

E. Estimation of the Noise Variance

The estimation of model parameters can be naturally implemented in SSoBaP by procedure (29), (30) described in Section IV-A. Considering a set of unknown parameters $\boldsymbol{\alpha}$, one can include $q(\boldsymbol{\alpha})$ as a new factor within the VB-EM equations and possibly add the additional constraint

$$q(\boldsymbol{\alpha}) = \delta(\boldsymbol{\alpha} - \hat{\boldsymbol{\alpha}}). \quad (47)$$

In the sequel, we will however not consider the *general* problem of model-parameter estimation, which can be particularly involved in Boltzmann machine. A lot of literature has already been dedicated to this problem and is out of the scope of this paper. We refer the interested reader to e.g., [52], [54], [67]–[69].

In this section, we exclusively focus on the estimation of the noise variance σ_n^2 which has revealed to be crucial for the algorithm performance in our empirical experiences. The noise variance can be seen as a disparity measure between the observation \mathbf{y} and its sparse approximation. Even if it is known *a priori*, its estimation turns out to be of great interest for the algorithm convergence. Indeed, SSoBaP relies on a successive refinement of the approximations of the posterior distributions $p(x_i, s_i|\mathbf{y})$'s, i.e., the sparse approximation: in the first iterations, the estimations are likely to be coarse, thus the disparity between \mathbf{y} and its sparse approximation might be large. The estimation of σ_n^2 at each iteration allows to take this evolution into account in the approximation process.

In practice, particularized to model (5), (6)–(10), we consider σ_n^2 as a new unknown variable in $\boldsymbol{\theta}$:

$$\boldsymbol{\theta} = [x_1 \ s_1 \ \dots \ x_M \ s_M \ \sigma_n^2]^T, \quad (48)$$

and add a factor in the MF structure (49) as

$$q(\mathbf{x}, \mathbf{s}, \sigma_n^2) = q(\sigma_n^2) \prod_i q(x_i, s_i). \quad (49)$$

Then, $q(\sigma_n^2)$ is constrained to

$$q(\sigma_n^2) = \delta(\sigma_n^2 - \hat{\sigma}_n^2), \quad (50)$$

leading to maximization (30), which becomes

$$\hat{\sigma}_n^2 = \arg \max_{\sigma_n^2} \left\{ \sum_{\mathbf{s}} \int \prod_i q(x_i, s_i) \log p(\mathbf{x}, \mathbf{s}, \mathbf{y}|\sigma_n^2) dx \right\}, \quad (51)$$

$$\begin{aligned}
&= \frac{1}{N} \left\langle \left\| \mathbf{y} - \sum_i s_i x_i \mathbf{d}_i \right\|^2 \right\rangle_{\prod_i q(x_i, s_i)} \quad (52) \\
&= \frac{1}{N} \left(\mathbf{y}^T \mathbf{y} - 2 \sum_i q(s_i = 1) m(s_i = 1) \mathbf{y}^T \mathbf{d}_i \right. \\
&\quad + \sum_i \sum_{j \neq i} q(s_i = 1) q(s_j = 1) m(s_i = 1) m(s_j = 1) \mathbf{d}_i^T \mathbf{d}_j \\
&\quad \left. + \sum_i q(s_i = 1) (\Sigma(s_i = 1) + m(s_i = 1)^2) \mathbf{d}_i^T \mathbf{d}_i \right). \quad (53)
\end{aligned}$$

Update (53) is inserted in procedure (36)–(41) after the estimation of the $q(x_i, s_i)$'s.

F. Relation to Past Work

In this subsection, we place the SSoBaP algorithm and its particular “unstructured” case, SoBaP, within the previous contributions of the literature. To be as exhaustive as possible, we identify the contributions considering Boltzmann machines, but also those using BG models, in a SR context:

1) *Boltzmann Machine*: The proposed SSoBaP can be compared to the three main contributions [52]–[54], which consider Boltzmann machines in a structured SR point of view. They mainly distinguish by the estimation problem they consider and the practical procedure they propose to solve it.

In [52], the authors focus on the MAP estimation of the support of the sparse representation (18) and propose a solution using Gibbs sampling and simulated annealing. The same estimation problem is considered by Peleg *et al.* in [54]. Emphasizing the high computational cost of the approach [52], they suggest a greedy alternative. The greedy approach is also adopted in [53] but to solve the joint MAP estimation problem (17). In this contribution, the authors derive the so-called LaMP (for “Lattice Matching Pursuit”), a structured version of CoSaMP.

In next Section V, we compare the proposed algorithm to the contributions [54], which presents a reasonable computational cost.

2) *BG Model*: As mentioned in the introduction, BG model (6), (7) has already been considered in some contributions [28], [29], [32], [33] and under the marginal formulation (3) in [35], [36]. However, all these contributions differ from the proposed approach by the estimation problem and the practical procedure introduced to solve it.

Thus, in [28], [29], [32], [33], the authors focus on the joint MAP estimation problem (17). They then propose different greedy procedures to solve it, some of them are explicitly related to standard deterministic algorithms, as BMP, BOMP [32] or SBR [33] and their respective standard homologues MP, OMP and OLS.

Contribution [50] considers a tree-structured version of BG model (6), (7) dedicated to a specific application (namely, the sparse decomposition of an image in wavelet or DCT bases). Besides this specific application, their approach relies, as ours, on a VB-EM algorithm. However, it differs by the MAP estimation problem (18) they address and the different MF factorization they choose to solve it. Finally, Ge *et al.* suggest in [34]

another approximation of $p(\mathbf{x}, \mathbf{s}|\mathbf{y})$ based on a MCMC inference scheme.

In contributions [35], [36], the authors use the marginal formulation (3). They propose to resort to the approximate message passing algorithm introduced in [44] and generalized in [70] to compute the posterior distribution of the sought sparse vector. Both also consider the possibility of estimating the parameters of the model (e.g., the noise variance, the Bernoulli parameter, the variance of the Gaussian distribution, etc.) by means of an Expectation-Maximization-like algorithm.

V. EXPERIMENTS

In this section, we study the performance of the proposed algorithm by extensive computer simulations. We assess the performance in terms of the reconstruction of the SR support and the estimation of the non-zero coefficients. To that end, we evaluate different figures of merit as a function of the number of atoms used to generate the data, say K . In particular, we consider empirical measures of the mean square error (MSE), the probability of missed detections, the probability of false detections. These figures are evaluated from 500 trials for each simulation points.

We assess the performance of the proposed algorithm in both the unstructured and structured cases and compare the results to those obtained with state-of-the-art procedures.

A. Unstructured Case

The unstructured case does not consider the possible structures existing between the atoms building the sparse representation. We use the following parameters: $N = 128$, $M = 256$, $\sigma_n^2 = 10^{-3}$ and generate the data as follows. Each point of simulation corresponds to a fixed number of non-zero coefficients K and, given this number, the positions of the non-zero coefficients are drawn uniformly at random for *each* observation. The elements of the dictionary are generated for each observation as realizations of a zero-mean Gaussian distribution with variance N^{-1} . The value of the non-zero coefficients in \mathbf{x} are generated according to the two different scenarios that we describe below.

We evaluate and compare the performance of 7 different algorithms: MP [16], SP [20], IHT [14], BP [10], BMP [32], EMBGAMP [36] and SoBaP. For SP, IHT, BP and EMBGAMP, we use the implementations available on author's webpages (resp. at <http://sites.google.com/site/igorcaron2/cscodes/>, <http://www.personal.soton.ac.uk/tb1m08/sparsify/sparsify.html>, <http://www.acm.caltech.edu/11magic/> (ℓ_1 -magic) and <http://www2.ece.ohio-state.edu/vilaj/EMBGAMP/EMBGAMP.html>). MP is run until the ℓ_2 -norm of the residual drops below $\sqrt{N\sigma_n^2}$. The same criterion is used for BP. BMP iterates as long as $\log p(\mathbf{y}, \hat{\mathbf{x}}^{(n)}, \hat{\mathbf{s}}^{(n)}) > \log p(\mathbf{y}, \hat{\mathbf{x}}^{(n-1)}, \hat{\mathbf{s}}^{(n-1)})$ where n is the iteration number (see [32]). SoBaP is run until $\forall i \in \{1, \dots, M\}, |q(s_i^{(n)}) - q(s_i^{(n-1)})| < 10^{-2}$. Finally, we set $p_i = \frac{K}{M}, \forall i$.

1) *Gaussian Model*: In this scenario, the amplitudes of the non-zero coefficients are drawn from a Gaussian distribution, according to (6). We set $\sigma_{x_i}^2 = 1\forall i$.

Fig. 1(a) shows the MSE as a function of the number of non-zero coefficients, K . For $K \leq 40$, SoBaP presents, together with EMBGAMP, the best performance. Beyond this bound, it




Fig. 1. MSE (a), probability of missed (b) and false (c) detection versus K . The support of the sparse vector is drawn uniformly at random. The non-zero coefficients in \mathbf{x} follow the “Gaussian model” and $\sigma_n^2 = 10^{-3}$.

is dominated by EMBGAMP but keeps a good behavior with regard to other algorithms.

Fig. 1(b) and (c) represents the algorithm performance in terms of the reconstruction of the SR support. We can observe that SoBaP succeeds in keeping both reasonable missed detection and false detection rates on a large range of sparsity levels. This is not the case for the other algorithms. If some of them




Fig. 2. Average running time versus K . The support of the sparse vector is drawn uniformly at random. The non-zero coefficients in \mathbf{x} follow the “Gaussian model” and $\sigma_n^2 = 10^{-3}$.

present better performance for one rate, this is at the expense of the other one. BP and EMBGAMP constitute two extreme examples. They are both based on a “spender” strategy: they prefer missing no atom [their missed detection is equal to zero, that is why they do not appear in Fig. 1(b)] even if they are not sure they are good ones.

It is difficult to compare the running times of the considered algorithms since they do not have the same stopping criteria. In Fig. 2, we see that SoBaP presents a computational cost higher than MP or BMP, while sharing a similar complexity order per update step (see Section IV-B). This can be explained by the fact that SoBaP updates all indices at each iteration, as we previously mentioned. Beyond these observations, SoBaP remains competitive with the other algorithms, in particular for high sparsities (i.e., small numbers of non-zero coefficients). Note finally that EMBGAMP constitutes here the most costly procedure, with a high constant running time.

2) “0–1” Model: In this second scenario, the amplitude of the non-zero coefficients in \mathbf{x} are forced to be equal to 1.

Fig. 3(a) shows the MSE as a function of the number of non-zero coefficients, K . For this particular setup, we experimentally observed that SoBaP presents better results when we set $\sigma_{x_i}^2 = 0.1\forall i$ in the algorithm. Thus, SoBaP outperforms all algorithms (or present similar performance, for high sparsities) except EMBGAMP which clearly dominates.

The performance achieved in terms of reconstruction of the SR support (see Fig. 3(b) and (c)) is similar to the one observed in the previous scenario. SoBaP constitutes the best compromise between missed and false detection rates, while BP and EMBGAMP follow the same strategy as before: they select all atoms, including “bad” ones (i.e., not used to generate the data).

B. Structured Case

In the structured case, the links between atoms are taken into account in the sparse decomposition. For the experiments, we considered two different structures: a Markov chain, for which we showed the equivalence with a particular Boltzmann machine in (13), (14), and a general, non-dedicated Boltzmann machine.




Fig. 3. MSE (a), probability of missed (b) and false (c) detection versus K . The support of the sparse vector is drawn uniformly at random. The non-zero coefficients in \mathbf{x} follow the “0–1” model and $\sigma_n^2 = 10^{-3}$.

1) *Markov Chain*: We first consider the simple scenario where the positions of the non-zero coefficients in \mathbf{x} follow a Markov chain model. We use the following parameters: $N = 128$, $M = 256$, $\sigma_n^2 = 10^{-3}$. The observations are generated as follows. The elements of the dictionary are drawn, for each observation, from a zero-mean Gaussian distribution with

variance N^{-1} . For each point of simulation, we fix the number of non-zero coefficients and select their positions uniformly at random. Either the Gaussian or the “0–1” model is considered for the amplitudes of the non-zero coefficients in \mathbf{x} .

We consider the case of a symmetric Markov chain i.e., $p_{i+1}^0 = p_{i+1}^1 \forall i$. The value of the p_{i+1}^0 's are drawn as follows

$$p_{i+1}^0 \sim \mathcal{U}[0, 0.5] \quad \text{if } s_i = s_{i+1}, \quad (54)$$

$$p_{i+1}^0 \sim \mathcal{U}[0.5, 1] \quad \text{otherwise.} \quad (55)$$

Boltzmann parameters \mathbf{b} and \mathbf{W} are then constructed according to (13), (14). The performance of SSoBaP is represented in Fig. 4 for both the Gaussian (dashed curves) and “0–1” (solid curves) models. SSoBaP is compared to the greedy procedure proposed by Peleg *et al.* in [54], called MAP—OMP-like. The latter also relies on a Boltzmann machine. The same parameters are thus used in both algorithms. The unstructured variant of SSoBaP, SoBaP, is considered too, in order to assess the relevance of accounting sparse structures with the BM parameters. MAP—OMP-like iterates until the ℓ_2 -norm of the residual drops below 10^{-3} or the iteration number exceeds $\frac{N}{2}$. SoBaP and SSoBaP are run until $|q(s_i^{(n)}) - q(s_i^{(n-1)})| < 10^{-2} \forall i \in \{1, \dots, M\}$.

We see in Fig. 4 that SSoBaP nicely takes benefit from the additional information on the SR support and thus improves the performance of SoBaP with respect to all figures of merit. In the Gaussian case, the probability of missed detection is roughly equal to $5 \cdot 10^{-2}$ (versus 10^{-1} for SoBaP) over a large range of sparsity levels with a probability of false detection of about 10^{-3} . In the “0–1” model, no missed detections have been detected up to $K = 58$ and the probability of false detection is of the order of 10^{-3} in this range. These good properties in terms of support recovery are confirmed in Fig. 4(a) by the MSE performance. Let us note, that for this particular scenario, MAP—OMP-like exhibits the worst performance. In particular, its probability of false detection rapidly increases as the number of non-zero coefficients becomes larger.

Finally, Fig. 5 illustrates the running time of the three procedures. We note that the computational burden induced by SoBaP strongly depends on the considered scenario. On the other hand, the running times of SSoBaP and MAP—OMP-like remain similar for both the Gaussian and “0–1” models. As far as this simulation setup is concerned, SSoBaP is significantly faster than MAP—OMP-like for small to moderate values of K .

2) *“General” Boltzmann Machine*: We consider the following parameters: $N = 32$, $M = 64$, $\sigma_n^2 = 10^{-3}$ and generate the data as follows. The elements of the dictionary are drawn, for each observation, from a zero-mean Gaussian with variance N^{-1} . For each point of simulation, we fix the number of non-zero coefficients *and* their positions in the SR support. These positions are thus drawn uniformly at random *once* for *all* observations. This leads to a particular support \mathbf{s} that we use for all trials. So, we average the performance of the algorithms on data structured in the same way. Regarding the amplitudes of the non-zero coefficients in \mathbf{x} , we consider the same scenarios as for the unstructured case, i.e., the Gaussian and “0–1” models.

The parameters of the Boltzmann machine, \mathbf{b} and \mathbf{W} , are drawn from the a posteriori distribution $p(\mathbf{b}, \mathbf{W} | \mathbf{s})$ by means of




Fig. 4. MSE (a), probability of missed (b) and false (c) detection versus K . The support of the sparse vector follows a Markov chain model. The non-zero coefficients in \mathbf{x} follow the “0–1” (solid) or “Gaussian” (dashed) model and $\sigma_n^2 = 10^{-3}$.

the “Single-variable Exchange” algorithm introduced in [67], using $w_{ij} \sim \mathcal{U}[-1, 1] \forall i, j$ and $b_i \sim \mathcal{U}[-20, 20] \forall i$ as a priori distributions. We initialize all elements in \mathbf{b} and \mathbf{W} to 0. For each point of simulation, the “Single-variable Exchange” algorithm is run with a burn-out iteration number of 1000; we then




Fig. 5. Average running time versus K . The support of the sparse vector follows the Markov chain model. The non-zero coefficients in \mathbf{x} follow the “0–1” (solid) or “Gaussian” (dashed) model and $\sigma_n^2 = 10^{-3}$.

allocate the 500 following parameter realizations to the 500 observations of the considered point.

SSoBaP is compared to MAP—OMP-like and SoBaP. MAP—OMP-like iterates until the ℓ_2 -norm of the residual drops below 10^{-3} or the iteration number exceeds $\frac{N}{2}$. SoBaP and SSoBaP are run until $|q(s_i^{(n)}) - q(s_i^{(n-1)})| < 10^{-2} \forall i \in \{1, \dots, M\}$.

Fig. 6(a), (b) and (c) sums up the performance achieved by the three algorithms under the two considered scenarios.

Focusing on the Gaussian model (dashed curves), we observe that SSoBaP dominates SoBaP and MAP—OMP-like in terms of MSE for a wide range of sparsity levels. Moreover, it presents stable missed and false detection rates (around 10^{-2}). We then can see that it outperforms SoBaP and MAP—OMP-like in terms of missed detection rate for all considered sparsities while achieving the lowest false detection rate for small sparsities ($K > 10$).

SSoBaP keeps its general good behavior with the “0–1” model (solid curves). This good behavior is even reinforced by zero missed detection for $K < 13$. Note that this does not contradict the similarity observed between the MSE curves: missed and false detection rates impact on the MSE but their influence is difficult to measure, as a high MSE can be due to high missed and false detection rates but also to a bad coefficients’ estimation.

Fig. 7 shows the running times of the considered algorithms in both the Gaussian and “0–1” scenarios. As far as these setups are concerned, SSoBaP has always a smaller running time than MAP—OMP-like. The behavior of SoBaP differs according to the considered scenario. For the Gaussian model (dashed curves), SoBaP has the smallest running time among the three algorithms. For the “0–1” model (solid curves), SoBaP is outperformed by SSoBaP.

VI. CONCLUSION

In this paper, we address the structured SR problem from a Bayesian point of view. Structures are taken into account by means of a Boltzmann machine which allows for the description of a large set of structures. We then focus on the resolution of




Fig. 6. MSE (a), probability of missed (b) and false (c) detection versus K . The support of the sparse vector follows the general model of a Boltzmann machine. The non-zero coefficients in \mathbf{x} follow the “0–1” (solid) or “Gaussian” (dashed) model and $\sigma_n^2 = 10^{-3}$.

marginalized MAP problems. The proposed approach is based on a mean-field approximation and the use of the “variational Bayes Expectation-Maximization” algorithm, and results in the so-called “Structured Soft Bayesian Pursuit” (SSoBaP) algorithm. We assess the performance of SSoBaP in the unstructured and structured cases (the unstructured version of SSoBaP is then




Fig. 7. Average running time versus K . The support of the sparse vector follows the general model of a Boltzmann machine. The non-zero coefficients in \mathbf{x} follow the “0–1” (solid) or “Gaussian” (dashed) model and $\sigma_n^2 = 10^{-3}$.

called SoBaP). In both cases, we evaluate the ability of the algorithm to reconstruct the SR support and estimate the non-zero coefficients. Experimental results show that the corresponding algorithms perform well in comparison to other state-of-the-art algorithms, at a reasonable computational cost.

Future work will consider the use of the proposed algorithm in practical applications, in particular in audio processing where structured sparsity can be favorably exploited for efficient representations of audio signals.

ACKNOWLEDGMENT

The authors wish to thank M. T. Faktor and Prof. M. Elad for providing their implementation of the MAP—OMP-like algorithm.

REFERENCES

- [1] L. Daudet, “Sparse and structured decompositions of audio signals in overcomplete spaces,” in *Proc. Int. Conf. Digit. Audio Effects (DAFx)*, Naples, Italy, Oct. 2004, pp. 22–26.
- [2] C. Févotte, L. Daudet, S. J. Godsill, and B. Torrèsani, “Sparse regression with structured priors: Application to audio denoising,” *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 16, no. 1, pp. 174–185, 2008.
- [3] B. D. Jeffs and M. Gunsay, “Restoration of blurred star field images by maximally sparse optimization,” *IEEE Trans. Image Process.*, vol. 2, no. 2, pp. 202–211, Apr. 1993.
- [4] R. M. Figueras i Ventura, P. Vanderghyest, and P. Frossard, “Low rate and scalable image coding with redundant representations,” EPFL, Lausanne, Switzerland, Tech. Rep., TR-ITS-03.02, Jun. 2003.
- [5] D. L. Donoho, “Compressed sensing,” *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [6] B. K. Natarajan, “Sparse approximate solutions to linear systems,” *SIAM J. Comput.*, vol. 24, no. 2, pp. 227–234, Apr. 1995.
- [7] A. Miller, *Subset Selection in Regression*, 2nd ed. London, U.K.: Chapman & Hall/CRC, Apr. 2002.
- [8] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani, “Least angle regression,” *Ann. Stat.*, vol. 32, no. 2, p. 407499, 2004.
- [9] H. Markowitz, “The optimization of a quadratic function subject to linear constraints,” *Nav. Res. Logistics Quarter.*, vol. 3, no. 1–12, pp. 111–133, Mar.–Jun. 1956.
- [10] S. S. Chen, D. L. Donoho, and M. A. Saunders, “Atomic decomposition by basis pursuit,” *SIAM J. Sci. Comput.*, vol. 20, pp. 33–61, 1998.
- [11] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *J. Roy. Stat. Soc.*, vol. 58, pp. 267–288, 1996.
- [12] I. F. Gorodnitsky and B. D. Bhaskar, “Sparse signal reconstruction from limited data using focus: A re-weighted minimum norm algorithm,” *IEEE Trans. Signal Process.*, vol. 45, no. 3, pp. 600–616, Mar. 1997.

- [13] N. G. Kingsbury and T. H. Reeves, "Overcomplete image coding using iterative projection-based noise shaping," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, 2002, vol. 3, pp. 597–600.
- [14] T. Blumensath and M. E. Davies, "Iterative thresholding for sparse approximations," *J. Fourier Anal. Appl.*, vol. 14, no. 5–6, pp. 629–654, Dec. 2008.
- [15] I. Daubechies, M. Defrise, and C. DeMol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Commun. Pure Appl. Math.*, vol. 11, pp. 1413–1457, 2004.
- [16] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3397–3415, Dec. 1993.
- [17] Y. C. Pati, R. Rezaifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, 1993, pp. 40–44.
- [18] C.-T. Chen, "Adaptive transform coding via quadtree-based variable blocksize DCT," presented at the IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP), May 23–26, 1989.
- [19] D. L. Donoho, Y. Tsaig, I. Drori, and J.-L. Starck, "Sparse solution of underdetermined linear equations by stagewise orthogonal matching pursuit," Stanford Univ., Stanford, CA, Tech. Rep., Mar. 2006.
- [20] W. Dai and O. Milenkovic, "Subspace pursuit for compressive sensing signal reconstruction," *IEEE Trans. Inf. Theory*, vol. 55, no. 5, pp. 2230–2249, May 2009.
- [21] D. Needell and J. A. Tropp, "Cosamp: Iterative signal recovery from incomplete and inaccurate samples," *Appl. Comput. Harmon. Anal.*, vol. 26, no. 3, pp. 301–321, May 2009.
- [22] B. A. Olshausen and D. J. Field, "Sparse coding with an overcomplete basis set: A strategy employed by v1?," *Vis. Res.*, vol. 37, no. 23, pp. 3311–3325, 1997.
- [23] M. S. Lewicki and T. J. Sejnowski, "Learning overcomplete representations," *J. Neural Comput.*, vol. 12, pp. 337–365, 2000.
- [24] M. Girolami, "A variational method for learning sparse and overcomplete representation," *J. Neural Comput.*, vol. 13, no. 11, pp. 2517–2532, 2003.
- [25] C. Févotte and S. J. Godsill, "A Bayesian approach for blind separation of sparse sources," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 14, no. 6, pp. 2174–2188, Nov. 2006.
- [26] C. Févotte and S. J. Godsill, "Blind separation of sparse sources using Jeffrey's inverse prior and the expectation-maximization algorithm," in *Proc. Int. Conf. Independent Component Anal. Blind Source Separation (ICA)*, 2006, pp. 593–600.
- [27] P. Schniter, L. C. Potter, and J. Ziniel, "Fast Bayesian matching pursuit," in *Proc. Workshop Inf. Theory Appl. (ITA)*, La Jolla, CA, Jan. 2008, pp. 326–333.
- [28] H. Zayyani, M. Babaie-Zadeh, and C. Jutten, "Bayesian pursuit algorithm for sparse representation," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, 2009, pp. 1549–1552.
- [29] H. Zayyani, M. Babaie-Zadeh, and C. Jutten, "An iterative bayesian algorithm for sparse component analysis in presence of noise," *IEEE Trans. Signal Process.*, vol. 57, no. 11, pp. 4378–4390, Nov. 2009.
- [30] D. Baron, S. Sarvotham, and R. G. Baraniuk, "Bayesian compressive sensing via belief propagation," Tech. Rep., Jun. 2009 [Online]. Available: <http://arxiv.org/0812.4627v2.pdf>
- [31] C. Herzet and A. Drémeau, "Sparse representation algorithms based on mean-field approximations," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, Dallas, TX, Mar. 2010, pp. 2034–2037.
- [32] C. Herzet and A. Drémeau, "Bayesian pursuit algorithms," presented at the Eur. Signal Process. Conf. (EUSIPCO), Aalborg, Denmark, Aug. 2010.
- [33] C. Soussen, J. Idier, D. Brie, and J. Duan, "From Bernoulli–Gaussian deconvolution to sparse signal restoration," CRAN/IRCCyN, Nantes, France, Tech. Rep., Jan. 2010.
- [34] D. Ge, J. Idier, and E. Le Carpentier, "Enhanced sampling schemes for MCMC based blind Bernoulli–Gaussian deconvolution," *Signal Process.*, vol. 91, pp. 759–772, Apr. 2011.
- [35] F. Krzakala, M. Mézart, F. Sausset, Y. F. Sun, and L. Zdeborová, "Statistical physics-based reconstruction in compressed sensing," [Online]. Available: <http://arxiv.org/abs/1109.4424v2> 2011
- [36] J. Vila and P. Schniter, "Expectation-maximization Bernoulli-Gaussian approximate message passing," presented at the Asilomar Conf. Signals, Syst., Comput., Monterey, CA, Nov. 2011.
- [37] Y. C. Eldar and M. Mishali, "Robust recovery of signals from a structured union of subspaces," *IEEE Trans. Inf. Theory*, vol. 55, no. 11, pp. 5302–5316, Nov. 2009.
- [38] Y. C. Eldar, P. Kuppinger, and H. Bolcskei, "Block-sparse signals: Uncertainty relations and efficient recovery," *IEEE Trans. Signal Process.*, vol. 58, no. 6, pp. 3042–3054, 2010.
- [39] M. Yuan and Y. Lin, "Model selection and estimation in regression with grouped variables," *J. Roy. Stat. Soc., Series B*, vol. 68, pp. 49–67, 2006.
- [40] L. Yu, J.-P. Barbot, G. Zheng, and H. Sun, "Compressive sensing for cluster structured sparse signals: Variational Bayes approach," Wuhan Univ.-ENSEA, Cergy-Pontoise, France, Tech. Rep., 2012.
- [41] L. Yu, H. Sun, J.-P. Barbot, and G. Zheng, "Bayesian compressive sensing for clustered sparse signals," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, Prague, Czech Republic, May 2011, pp. 3948–3951.
- [42] J. Huang, T. Zhang, and D. Metaxas, "Learning with structured sparsity," in *Proc. Int. Conf. Mach. Learn.*, 2009, pp. 1–8.
- [43] S. Rangan, A. K. Fletcher, V. K. Goyal, and P. Schniter, "Hybrid approximate message passing with applications to structured sparsity," [Online]. Available: <http://arxiv.org/abs/1111.2581> 2011
- [44] D. L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing: I. motivation and construction," in *Proc. IEEE Inf. Theory Workshop (ITW)*, Cairo, Egypt, Jan. 2010, pp. 1–5.
- [45] P. Sprechmann, I. Ramirez, G. Sapiro, and Y. Eldar, "Collaborative hierarchical sparse modeling," in *Proc. IEEE Int. Conf. Inf. Sci. Syst. (CISS)*, Mar. 2010, pp. 1–6.
- [46] M. Kowalski and B. Torrèrani, "Sparsity and persistence: Mixed norms provide simple signal models with dependent coefficients," *Signal, Image, Video Process.*, vol. 3, no. 3, pp. 251–264, 2009.
- [47] L. Daudet, "Sparse and structured decompositions of signals with the molecular matching pursuit," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 14, no. 5, pp. 1808–1816, Sep. 2006.
- [48] R. Jenatton, J. Mairal, G. Obozinski, and F. Bach, "Proximal methods for hierarchical sparse coding," INRIA, Rocquencourt, France, Tech. Rep., 2010.
- [49] L. He and L. Carin, "Exploiting structure in wavelet-based Bayesian compressive sensing," *IEEE Trans. Signal Process.*, vol. 57, no. 9, pp. 3488–3497, Sep. 2009.
- [50] L. He, H. Chen, and L. Carin, "Tree-structured compressive sensing with variational bayesian analysis," *IEEE Signal Process. Lett.*, vol. 17, no. 3, pp. 233–236, 2010.
- [51] P. Schniter, "Turbo reconstruction of structured sparse signals," in *Proc. IEEE Annu. Conf. Inf. Sci. Syst. (CISS)*, Princeton, NJ, Mar. 2010, pp. 1–6.
- [52] P. J. Garrigues and B. A. Olshausen, "Learning horizontal connections in a sparse coding model of natural images," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, Dec. 2008, pp. 505–512.
- [53] V. Cevher, M. F. Duarte, C. Hegde, and R. G. Baraniuk, "Sparse signal recovery using Markov random fields," presented at the Adv. Neural Inf. Process. Syst. (NIPS), Vancouver, Canada, Dec. 2008.
- [54] T. Peleg, Y. C. Eldar, and M. Elad, "Exploiting statistical dependencies in sparse representations for signal recovery," *IEEE Trans. Signal Process.*, vol. 60, no. 5, pp. 2286–2303, May 2012.
- [55] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hedge, "Model-based compressive sensing," *IEEE Trans. Inf. Theory*, vol. 56, pp. 1982–2001, Apr. 2010.
- [56] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, "A learning algorithm for Boltzmann machines," *Cogn. Sci.*, vol. 9, no. 1, pp. 147–169, 1985.
- [57] B. C. Levy, *Principles of Signal Detection and Parameter Estimation*, 1 ed. New York: Springer, Jul. 2008.
- [58] M. J. Wainwright and M. I. Jordan, "Graphical models, variational inference and exponential families," Dept. of Statistics, Univ. of California, Berkeley, CA, Tech. Rep., 2003.
- [59] M. Beal, "Variational Algorithms for Approximate Bayesian Inference," Ph.D. thesis, Univ. College of London, London, U.K., May 2003.
- [60] T. P. Minka, "Using lower bounds to approximate integrals," Media Lab, Mass. Inst. of Technol., Cambridge, MA, Jun. 2001.
- [61] M. J. Beal and Z. Ghahramani, "The variational bayesian em algorithm for incomplete data: With application to scoring graphical model structures," *Bayesian Stat.*, vol. 7, pp. 453–463, 2003.
- [62] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York: Springer, 2006.
- [63] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the em algorithm," *J. Roy. Stat. Soc., Series B (Methodologic.)*, vol. 39, pp. 1–38, 1977.
- [64] R. M. Neal and G. E. Hinton, "A view of the em algorithm that justifies incremental, sparse, and other variants," *Learn. Graph. Models*, vol. 89, pp. 355–368, 1998.
- [65] A. Drémeau, C. Herzet, and L. Daudet, "Soft Bayesian pursuit algorithm for sparse representations," in *Proc. IEEE Int. Stat. Signal Process. Workshop (SSP)*, 2011, pp. 341–344.

- [66] H. Zayyani, M. Babaie-Zadeh, and C. Jutten, "Sparse component analysis in presence of noise using em-map," in *Proc. Int. Conf. Independent Component Anal. Signal Separation*, London, 2007.
- [67] I. Murray, Z. Ghahramani, and D. J. C. MacKay, "Mcmc for doublyintractable distributions," in *Proc. Annu. Conf. Uncertainty in Art. Intell. (UAI)*, 2006, pp. 359–366, AUA Press.
- [68] M. J. Nijman and H. J. Kappen, "Efficient learning in sparsely connected Boltzmann machines," in *Proc. Int. Conf. Art. Neural Netw.*, 1996, pp. 41–46.
- [69] N. L. Lawrence, C. M. Bishop, and M. I. Jordan, "Mixture representations for inference and learning in Boltzmann machines," in *Proc. Conf. Uncertainty Art. Intell.*, 1998, pp. 320–327.
- [70] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," [Online]. Available: <http://arxiv.org/abs/1010.5141> 2010



Angélique Drémeau received the State Engineering degree from Télécom Bretagne, Brest, France, in 2007 and the M.Sc. and the Ph.D. degree in signal processing and telecommunications from the Université de Rennes, Rennes, France, in 2007 and 2010, respectively.

During her Ph.D. studies, she was with the Institut National de Recherche en Informatique et Automatique (INRIA), Rennes, France. In December 2010, she joined the Institut Langevin, ESPCI ParisTech, Paris, France, as a Postdoctoral Researcher, and since

September 2011, she has been a Postdoctoral Researcher at Télécom ParisTech, Paris, France.



Cédric Herzet received the Electrical Engineering degree and the Ph.D. degree in Applied Science from the Université Catholique de Louvain (UCL), Louvain-la-Neuve, Belgium, respectively, in 2001 and 2006. His Ph.D. thesis dealt with iterative synchronization algorithms for digital burst communications.

From August 2001 to April 2006, he was a Research Assistant in the Communications and Remote Sensing Laboratory at UCL. From May 2006 to December 2007, he was a Postdoctoral

Researcher within the Ecole normale supérieure de Cachan, Paris, France, and the University of California, Berkeley (under a Fulbright scholarship). He is currently a Researcher within the Institut National de Recherche en Informatique et Automatique (INRIA), Rennes, France. His topics of research include algorithm design for digital communications, image processing and compression, low-rank approximations and sparse representation algorithms.



Laurent Daudet (M'04–SM'10) studied at the Ecole Normale Supérieure in Paris, where he graduated in statistical and non-linear physics. In 2000, he received a PhD in mathematical modeling from the Université de Provence, Marseille, France. After a Marie Curie post-doctoral fellowship at the C4DM, Queen Mary University of London, UK, he worked as associate professor at UPMC (Paris 6 University) in the Musical Acoustics Lab. He is now Professor at Paris Diderot University—Paris 7, with research at the Langevin Institute for Waves and Images, where

he currently holds a joint position with the Institut Universitaire de France. Laurent Daudet serves as associate editor for the IEEE TRANSACTIONS ON AUDIO, SPEECH, AND LANGUAGE PROCESSING, and is author or co-author of over 100 publications (journal papers or conference proceedings) on various aspects of acoustics and audio signal processing, in particular using sparse representations.

Reconstruction of Dispersion Curves in the Frequency-Wavenumber Domain Using Compressed Sensing on a Random Array

Angélique Drémeau, Florent Le Courtois, and Julien Bonnel, *Member, IEEE*

Abstract—In underwater acoustics, shallow-water environments act as modal dispersive waveguides when considering low-frequency sources, and propagation can be described by modal theory. In this context, propagated signals are composed of few modal components, each of them propagating according to its own wavenumber. Frequency-wavenumber ($f - k$) representations are classical methods allowing modal separation. However, they require large horizontal line sensor arrays aligned with the source. In this paper, to reduce the number of sensors, a sparse model is proposed and combined with prior knowledge on the wavenumber physics. The method resorts to a state-of-the-art Bayesian algorithm exploiting a Bernoulli–Gaussian model. The latter, well suited to the sparse representations, makes possible a natural integration of prior information through a wise choice of the Bernoulli parameters. The performance of the method is quantified on simulated data and finally assessed through a successful application on real data.

Index Terms—Acoustic signal processing, compressed sensing, geoaoustic inversion.

I. INTRODUCTION

WHEN considering shallow-water zones and low-frequency sound propagation, the sound field is described by a sum of a few dispersive modes. In this context, matched mode processing [1] constitutes a relevant approach to infer the properties of the environment [2]. To be effective, this processing requires estimating the modes from the data with a great accuracy.

This modal estimation step has been studied considering various configurations. The most classical one is modal filtering using vertical line arrays (VLAs) [3], [4] but other receiver setups have also been addressed as horizontal line arrays (HLAs) or single hydrophones. In the case of single hydrophones, we can distinguish two different situations. If the source is motionless, then the modes can be separated using time-frequency processing when the range is large enough [5], and nonlinear signal processing when the range gets smaller [6], [7]. If the source is moving, then a synthetic horizontal aperture may be formed to

resolve mode interferences [8], [9] or modal wavenumbers [10], [11]. True horizontal aperture is directly provided by HLAs. Essentially, if the acoustic field can be sampled in the range (i.e., horizontal) dimension, then the mode estimation procedure is equivalent to a spectral estimation problem. As a result, considering a long HLA and a monochromatic source at the endfire position, the wavenumber spectrum can be obtained by applying a Fourier transform in the array dimension [12], [13]. In this paper, we propose a new method for wavenumber spectrum estimation in the context of short HLAs and broadband sources. Our work is particularized to shallow-water environments and low frequencies. In addition, we assume the environment to be range independent, and/or varying slowly so that mode coupling can be neglected.

In the following, the spatial Fourier transform will be denoted SFT. It transforms signals from the spatial to the wavenumber domain. The notation TFT will be reserved for the (more classical) temporal Fourier transform, which transforms signals from the time domain to the frequency domain. The representation associated to the wavenumber domain will always be referred to as the wavenumber spectrum, while the representation associated to the frequency domain will always be referred to as the spectrum.

Wavenumber spectrum estimation using SFT does not present any theoretical difficulty. However, it suffers from the classical drawbacks of the TFT. A large aperture is required to obtain a decent resolution and the hydrophone spacing must be small enough to prevent aliasing. Because of these two restrictions, extremely large HLAs are required to resolve the modes in underwater acoustics using SFT. However, this issue may be circumvented by resorting to more advanced spectral estimation methods. Among them, high-resolution (HR) methods are known to be more accurate than the SFT. When using the same number of measurements, HR methods allow a better separation of nearby spectral components [14]. Applications of HR methods in underwater acoustics for modal separation have been proposed using autoregressive models [15]–[18] and subspace separation methods [19], [20].

Since the propagation in shallow-water environments is described by a small number of modes, the consideration of sparse models constitutes a promising alternative to HR methods. Developments in signal processing have raised the interest of compressed sensing (CS) methods for an accurate sampling and reconstruction of sparse signals [21]. CS methods have met many

Manuscript received February 25, 2016; accepted December 20, 2016. Date of publication March 13, 2017; date of current version October 11, 2017. This work has been supported by DGA/MRIS. (*Corresponding author: Angélique Drémeau.*)

Associate Editor: N. Chotiros.

The authors are with ENSTA Bretagne and Lab-STICC (UMR CNRS 6285), 29806 Brest Cedex, France (e-mail: angelique.dremeau@ensta-bretagne.fr).

Digital Object Identifier 10.1109/JOE.2016.2644780




Fig. 1. Illustration of a $f - k$ plane obtained within configuration setup exposed in Section IV-A using 240 sensors.

applications in underwater acoustics [22], [23], and for imaging the modal dispersion of surface waves [24]. The first application of CS to estimate modal spectrum in underwater acoustics has been proposed in [25].

When considering a broadband source, wavenumber spectra can be estimated at each of the source frequencies. The concatenation of them results in a frequency-wavenumber ($f - k$) diagram (see Fig. 1), representative of the waveguide dispersion [12], [13], [25]. As a consequence, the $f - k$ diagram constructions suffer from the drawbacks inherited from each wavenumber spectrum estimation. To partially solve this problem, Le Courtois and Bonnel have proposed [26] a postprocessing method to track the wavenumbers in badly resolved $f - k$ diagrams. Tracking is performed using particle filtering (PF) [27]. To do so, the wavenumber spectrum (at a given frequency) is modeled as a dynamic system parametrized by two equations: a system equation and an observation equation. The observation equation is the equation that allows generating a wavenumber spectrum if the (discrete) wavenumbers are known. On the other hand, the system equation is an iterative relation linking the wavenumbers from one frequency to another. This system equation thus benefits from a physical hypothesis on the wavenumber dispersion in the frequency dimension, and such an idea will also be used in our paper. Note that the PF method is particularly interesting when the number of sensors is too small to separate the modes. However, the number of propagating modes must be known *a priori* to initiate the tracking operation. In addition, the resolution of the method depends on the distribution of the particles and cannot be easily estimated. These last two constraints will be eased in this paper.

This paper introduces a framework relying on the hypothesis that the propagation is sparse (in the wavenumber dimension) and dispersive (in the frequency dimension). A methodology allowing $f - k$ estimation and integrating these two physical hypotheses is presented. To this end, the work focuses on the sparse estimation of the wavenumber spectrum using a Bernoulli-Gaussian model [28]. In this framework, each wavenumber spectrum is a sparse vector. Each wavenumber bin is supposed to be “on” (i.e., there is a mode propagating at this particular

wavenumber value) or “off” (i.e., there is no mode propagating at this particular wavenumber value) according to a Bernoulli distribution. Once a wavenumber bin is “on,” the corresponding mode amplitude is supposed to follow a Gaussian distribution. In this approach, the Bernoulli parameter offers a great opportunity to take into account the dispersive propagation, so that the sparse wavenumber spectra estimated at each frequency are naturally related with each other. In this paper, the Soft Bayesian Pursuit (SoBaP) algorithm [29] is considered to efficiently perform the estimation procedure.

The paper is organized as follows. Section II recalls the physical principles of the modal propagation and the prior knowledge exploited to build up the $f - k$ representation. Section III is dedicated to a mathematical formulation of the problem and the integration of the physical prior. Finally, in Section IV, the method is applied to synthetic simulations and to marine data recorded in the North Sea.

II. ACOUSTIC PROPAGATION IN DISPERSIVE SHALLOW-WATER ENVIRONMENTS

A. Received Signal

In shallow-water environments, the acoustic propagation is described by modal theory. The propagation depends on the frequency f and is dispersive. When considering an emitting source $s(f)$ at depth z_s , the received signal on a sensor located at distance r and depth z can be written as [30]

$$y(f, r) = Q \frac{s(f)}{\sqrt{r}} \sum_{m=1}^{M(f)} \psi_m(f, z_s) \psi_m(f, z) \frac{e^{-jrk_{rm}(f)}}{\sqrt{k_{rm}(f)}} + w(f, r) \quad (1)$$

where Q is a constant factor, $M(f)$ is the number of propagating modes at frequency f , $k_{rm}(f)$ is the horizontal wavenumber of the m th mode, and $\psi_m(f, z)$ is the modal depth function of the m th mode. The quantity $w(f, r)$ stands for the TFFT of the noise attached to the measurements. Note that, for the sake of clarity, we deliberately omit here the dependence in z in the notation of $y(f, r)$ and $w(f, r)$ as we will consider HLAs, i.e., sensors at a constant depth.

For subsequent convenience, we define the amplitude of the m th mode as

$$A_m(f) \triangleq \frac{\psi_m(f, z_s) \psi_m(f, z)}{\sqrt{k_{rm}(f)}} \quad (2)$$

so that (1) can finally be reexpressed as

$$y(f, r) = Q \frac{s(f)}{\sqrt{r}} \sum_{m=1}^{M(f)} A_m(f) e^{-jrk_{rm}(f)} + w(f, r). \quad (3)$$

Considering an HLA of regularly spaced sensors aligned with the source (endfire position), the ℓ th sensor is at a distance of $r = (\ell - 1)\Delta_r + r_0$ from the source, where r_0 is the distance of the first sensor and Δ_r is sensor spacing. In this paper, we assume a distant source (so that $r_0 \gg \Delta_r$). Under this assumption, the geometrical attenuation factor $1/\sqrt{r}$ can be considered as constant over the entire sensor array. Similarly, we will neglect

any eventual modal attenuation, so that the modal wavenumber $k_{rm}(f)$ may be considered as real numbers.

As stated in the introduction, considering a monochromatic source, the wavenumber spectrum can be estimated by performing a simple SFT along the HLA. If a broadband source is available, an $f - k$ diagram can be obtained by concatenating the wavenumber spectra at several frequencies [12]. However, this requires a long and dense HLA. The goal of this paper is to estimate $f - k$ diagram in less constrained contexts, in particular, when the number of sensors is low.

B. Dispersion Relation

In waveguides, the horizontal wavenumbers k_{rm} are linked to their vertical counterparts k_{zm} by the dispersion relation, that is, for a given frequency f

$$\left(\frac{2\pi f}{c}\right)^2 = k_{rm}(f)^2 + k_{zm}(f)^2 \quad (4)$$

where c is the speed of sound.¹

Discretizing the frequency axis (with $f = \nu\Delta_f$, $\nu \in \mathbb{N}$) and denoting $k_{rm}[\nu] = k_{rm}(\nu\Delta_f)$, the wavenumbers attached to two successive indices are then defined as [25]

$$k_{rm}[\nu + 1]^2 = k_{rm}[\nu]^2 + (2\nu + 1) \left(\frac{2\pi\Delta_f}{c}\right)^2 + \epsilon[\nu] \quad (5)$$

where Δ_f is the sampling period, i.e., the step between two frequency bins, and $\epsilon[\nu] = k_{zm}[\nu]^2 - k_{zm}[\nu + 1]^2$. In shallow-water environments, the vertical wavenumbers k_{zm} weakly depend on the frequency [30]; the quantity ϵ is smaller than the other terms of the equation and can be neglected. This approximation may not be true when the propagation becomes strongly range dependent. This case will not be considered in the remainder of the paper.

Equation (5) expresses the dispersive nature of the wavenumbers in a shallow-water environment. In this paper, we propose to take into account this physical information in the reconstruction procedure of the $f - k$ diagram. Note that (5) was used in [26] and [25] to postprocess badly resolved $f - k$ diagrams. In this paper, (5) will be embedded as a prior information into a Bayesian CS algorithm, allowing a direct estimation of the $f - k$ diagram.

III. COMPRESSED SENSING

When dealing with compressed sensing, two properties have to be verified:

- 1) sparsity: the signal to be acquired can be represented with a few nonzero elements in a given representation basis;
- 2) incoherence: the sensing must be made in a domain “as orthogonal as possible” to the representation basis.

Under these conditions, the theory of compressed sensing guarantees that we can recover the signal from a number of

¹Note that quantities c and $k_{zm}(f)$ are depth dependent. Here again, we choose to skip this dependence in the expression as it will be considered as constant in our sensing context. As we will see in the following, the way the dispersion relation will be exploited in the proposed approach allows some deviations from this assumption.

samples of the order of the number of its nonzero elements in the representation basis.

In other words, “incoherence” expresses the fact that a signal admitting a sparse representation in a particular domain (for example, in the frequency domain) must be acquired in the domain where its representation is spread out (for example, the time domain). In the time-frequency example, the sensing domain is the canonical basis (time) and the sparsity domain is the Fourier basis (frequency).

The reconstruction problem of the $f - k$ diagram is then a perfect application case for compressed sensing: the signal of interest is sparse in the wavenumber domain (only few wavenumbers propagate in shallow-water environments and at low frequencies), and is acquired in the spatial domain. From a mathematical point of view, this means that we acquire the signal in the canonical basis (space, through the sensor array) and the sparsity domain is the spatial Fourier basis (wavenumber). We refer the reader to [31] for a brief introduction to compressed sensing, or to [32] for more insights.

A. State of the Art

Formally, let $\mathbf{y}_\nu \in \mathbb{C}^L$ be the signal measured over the L sensors at the frequency index ν . Adopting a discretized matrix formulation, (3) can be reexpressed as

$$\mathbf{y}_\nu = \mathbf{D}\mathbf{z}_\nu + \mathbf{w}_\nu \quad (6)$$

where \mathbf{D} is an $(L \times N)$ -dictionary of Fourier discrete atoms, namely whose (ℓ, n) -element is $d_{nl} \triangleq e^{-j2\pi\frac{n\ell}{N}}$, and $\mathbf{z}_\nu = [z_{\nu,1}, \dots, z_{\nu,N}]^T \in \mathbb{C}^N$ is the wavenumber spectrum at the frequency index, i.e., the ν th transposed line of the $f - k$ diagram to estimate.

Then, N corresponds to the number of discretized points in the horizontal wavenumber domain, say

$$\kappa_{rn} \triangleq 2\pi\frac{n}{N\Delta_r} \quad \forall n \in \{1, \dots, N\} \quad (7)$$

and \mathbf{z}_ν gathers the amplitudes attached to each of the wavenumber bins.

According to the modal theory in shallow-water environments and at low frequencies, we have $M[\nu] \ll N$ [see (3)], where $M[\nu] = M(\nu\Delta_f)$. In other words, the vector \mathbf{z}_ν has few nonzero elements, corresponding to the propagating modal wavenumbers. Formally, we define by $\mathcal{M}_\nu \subset \{1, \dots, N\}$ the set of cardinality $M[\nu]$ gathering the indices of the propagating wavenumbers. For each of the $M[\nu]$ propagating modes, there exists $n \in \mathcal{M}_\nu$ such as,

$$k_{rm}[\nu] = \kappa_{rn} \quad \text{and} \quad A_m[\nu] = z_{\nu,n} \quad (8)$$

where $A_m[\nu] = A_m(\nu\Delta_f)$ defined as in (2).

We note here that the precise estimation of the set \mathcal{M}_ν is crucial. From it derives the knowledge of the propagating wavenumbers $k_{rm}[\nu]$ and the modal amplitudes $A_m[\nu]$, of particular interest for source depth estimation and/or environmental inversion.

The sparsity of the vector \mathbf{z}_ν constitutes important information on the $f - k$ diagram, which should be taken into account

in the reconstruction procedure. Several formulations of the corresponding sparse recovery problem can then be considered. In this paper, we focus on the following one:

$$\hat{\mathbf{z}}_\nu = \underset{\mathbf{z}_\nu}{\operatorname{argmin}} \|\mathbf{y}_\nu - \mathbf{D}\mathbf{z}_\nu\|_2^2 + \lambda \|\mathbf{z}_\nu\|_0 \quad (9)$$

where $\|\mathbf{z}_\nu\|_0$ stands for the ℓ_0 -pseudonorm of \mathbf{z}_ν (counting the number of nonzero elements in \mathbf{z}_ν) and λ is a parameter specifying the tradeoff between the sparsity constraint and the data-fidelity term $\|\mathbf{y}_\nu - \mathbf{D}\mathbf{z}_\nu\|_2^2$. The solution of (9) can be naturally interpreted as the least square (LS) solution (which, in our case, is equivalent to an SFT when \mathbf{y}_ν is uniformly sampled) penalized by the sparsity of the solution.

Solving (9) is an NP-hard problem [33], i.e., it generally requires a combinatorial search over the entire solution space. Therefore, heuristic (but tractable) algorithms have been devised to deal with this problem. We can roughly divide them into three families: 1) the greedy algorithms (e.g., [34] and [35]) which build up the sparse solution by making a succession of greedy decisions; 2) the relaxation-based algorithms (e.g., [36] and [37]) which replace the ℓ_0 -pseudonorm by some ℓ_p -norm (with $p \in]0, 1[$) leading to a relaxed problem efficiently solvable by standard optimization procedures; and 3) the Bayesian algorithms (e.g., [38], [29], and [39]) which express the problem as the solution of a Bayesian inference problem and apply statistical tools to solve it.

The Bayesian approaches are particularly suitable when the noise level is not known. In addition, they offer a simple framework to implement any prior information available on the physical context of the problem. In our case, they present in particular a great opportunity to take into account the dispersion relation given in (5).

B. Bayesian Formulation

Let us first assume \mathbf{w}_ν in (6) to be a circular Gaussian noise (denoted by \mathcal{CN}) with zero mean and variance σ_w^2 . We suppose then that \mathbf{z}_ν is the realization of a Bernoulli–Gaussian (BG) model [29], that is

$$\mathbf{z}_\nu = \mathbf{s}_\nu \odot \mathbf{x}_\nu \quad (10)$$

where \odot represents here the term-by-term product, and

$$p(\mathbf{s}_\nu) = \prod_{n=1}^N p(s_{\nu,n}), \quad \text{with } p(s_{\nu,n}) = \operatorname{Ber}(p_{\nu,n}) \quad (11)$$

$$p(\mathbf{x}_\nu) = \prod_{n=1}^N p(x_{\nu,n}), \quad \text{with } p(x_{\nu,n}) = \mathcal{CN}(0, \sigma_x^2). \quad (12)$$

The Bernoulli distribution (denoted by Ber in the equation above) has a realization domain on $\{0, 1\}$ and depends on a parameter $p_{\nu,n}$ which represents the probability of being equal to 1. The circular Gaussian distribution put on the variable \mathbf{x}_ν is assumed to be with zero mean and variance σ_x^2 .

In other words, \mathbf{s}_ν is called the support of the sparse representation. It is such as $\mathbf{s}_\nu \in \{0, 1\}^N$. Every 1-value in \mathbf{s}_ν indicates a wavenumber bin corresponding to a propagating

modal wavenumber; every 0-value in \mathbf{s}_ν indicates a wavenumber bin corresponding to a wavenumber that does not propagate. The quantity \mathbf{x}_ν stands for the amplitude of sparse representation. Coming back to (8), \mathbf{s}_ν serves then to define the set \mathcal{M}_ν of the wavenumber bins corresponding to the propagating wavenumbers

$$\mathcal{M}_\nu = \{n \in \{1, \dots, N\} | s_{\nu,n} = 1\}$$

while \mathbf{x}_ν can be directly linked to the amplitudes $A_m[\nu]$ of the propagating modes. As a consequence, the latter are assumed to be independently and identically distributed according to a Gaussian law, as expressed in (12).

Formally, the BG model (11)–(12) is well suited to modeling situations where \mathbf{y}_ν stems from a sparse process: if $p_{\nu,n} \ll 1$, $\forall n$, only a small number of $s_{\nu,n}$'s will typically be nonzero, i.e., the observations \mathbf{y}_ν will be generated with high probability from a small subset of the columns of \mathbf{D} . More rigorously, it has been proved (see [28] and [40]) that, within model (11)–(12), the joint maximum *a posteriori* estimation problem shares the same set of solutions as the standard sparse recovery problem (9). This connection gives weight to the general use of this model in sparse recovery problems.

In this paper, we choose to resort to a particular algorithm of the sparsity literature, that is the SoBaP algorithm [29]. Exploiting model (6)–(12), SoBaP aims at solving the following marginalized MAP problem:

$$\hat{\mathbf{s}}_\nu = \underset{\mathbf{s}_\nu}{\operatorname{argmax}} p(\mathbf{s}_\nu | \mathbf{y}_\nu) \quad (13)$$

where

$$p(\mathbf{s}_\nu | \mathbf{y}_\nu) = \int_{\mathbf{x}_\nu} p(\mathbf{s}_\nu, \mathbf{x}_\nu | \mathbf{y}_\nu) d\mathbf{x}_\nu. \quad (14)$$

The algorithm relies on a variational approximation [41] which constitutes the building block of the current state-of-the-art algorithms [42].

Once \mathbf{s}_ν is estimated, the amplitudes of the propagated wavenumbers can be computed by a simple pseudoinversion of the dictionary \mathbf{D} , restricted to its nonzero columns, say $\mathbf{D}_{\hat{\mathbf{s}}_\nu}$

$$\hat{\mathbf{x}}_\nu = \mathbf{D}_{\hat{\mathbf{s}}_\nu}^+ \mathbf{y}_\nu \quad (15)$$

where $^+$ stands for the pseudoinverse operator.

C. Incorporating A Priori Information About Dispersive Propagation

In our application context, the BG model presents an additional advantage: relying on the dedicated support variable \mathbf{s}_ν , it allows for an easy implementation of *a priori* information through the Bernoulli parameters $p_{\nu,n}$ [see (11)].

By the physical relation (5), the wavenumbers estimated at the frequency ν define *a priori* values for the wavenumbers at the next frequency $\nu + 1$. Consequently, the sparse representation support can be propagated to $\nu + 1$. This can be done by means of the Bernoulli parameters $p_{\nu+1,n}$ which set the probabilities for the elements in $\mathbf{s}_{\nu+1}$ to be set to 1 (i.e., for the wavenumbers




Fig. 4. Reconstruction of the dispersion curves (in decibels) using measurements from 30 sensors and an SNR of 10 dB: (a) LS inversion; (b) OMP; and (c) SoBaP with dispersive *a priori*. The ground truth is represented by red points.




Fig. 5. Normalized MSE on the estimated wavenumbers (in decibels) as a function of: (a) the SNR for an array of $L = 120$ sensors; and (b) the number of sensors in the array for an SNR of 10 dB.

50 Hz (corresponding to a white noise or an impulsive signal in the time domain; the latter case was simulated here). The frequency resolution is $\Delta_f = 0.2$ Hz. The source is placed at $D = 130$ -m depth to avoid nodes of the modal functions. At 50 Hz, five modes are propagating.

Two sparse representation algorithms are compared for the resolution of the problem: orthogonal matching pursuit (OMP) [43] as used in [24], and SoBaP incorporating *a priori* information on the dispersive propagation, as defined earlier. OMP is stopped when $\|\mathbf{y}_\nu - \mathbf{D}\hat{\mathbf{x}}_\nu\|_2 < \sqrt{L\sigma_w^2}$. SoBaP is stopped when the Kullback–Leibler divergence between the distribution $p(\mathbf{s}, \mathbf{x}|\mathbf{y})$ and its variational approximation is smaller than 10^{-4} . Both algorithms are compared to the least square (LS) approach, performing a simple inverse Fourier transform.

The resulting $f - k$ diagrams are shown in Fig. 4 together with the ground truth, represented by red points. For all three, 30 sensors are used and a signal-to-noise ratio (SNR) of 10 dB is considered. To account for the CS framework, the positions of the sensors are randomly selected from the 240 simulated measurements. At these settings, the $f - k$ diagram obtained by simple LS inversion and plotted in Fig. 4(a) cannot be correctly interpreted to infer the expected wavenumbers: they do not appear clearly. The one obtained by OMP seems to be more

sensitive to noise; some wavenumbers are identified but they have no physical justification. The use of the dispersive *a priori* information in SoBaP suppresses all of the artifacts from the $f - k$ plane. These visual results are to be compared with the ground truth, represented by red points on each $f - k$ diagram.

To quantify more precisely the performance of the algorithms, we compute for each of them the normalized mean square error (MSE) between the estimated wavenumbers (corresponding, for each frequency index ν , to the $M[\nu]$ largest coefficients in $\hat{\mathbf{x}}_\nu$) and the true values of the Pekeris model. The error is averaged over the mode number and the frequencies to obtain a global value. The MSE is represented, in Fig. 5(a), as a function of the SNR for a given number of sensors $L = 120$ and in Fig. 5(b), as a function of the number of hydrophones exploited to reconstruct the $f - k$ diagram for a given SNR of 10 dB.

In accordance with previous works [25], [24], both figures illustrate the good behavior of sparse-aware algorithms with regard to naïve LS inversion: OMP and the proposed SoBaP procedure outperform the LS estimation in most practical setups, i.e., noisy measurements and few sensors.

As shown in Fig. 5(a), for an SNR greater than 20 dB, the two sparse-aware methods present similar performances, and an exact reconstruction of the wavenumbers is obtained (note




Fig. 6. Reconstruction of the dispersion curves (in decibels) from the North Sea measurements using: (a) 240 sensors with an LS inversion; (b) 20 sensors with an LS inversion; (c) 20 sensors with the OMP algorithm; and (d) 20 sensors with the physics-aware SoBaP algorithm. Red points represent the result of a peak-picking method on the LS inversion on the entire antenna (a).

that the residual MSE is linked to the resolution of the Fourier matrix \mathbf{D} , independently of the number of measurements). The proposed approach proves to be more interesting in the noisy cases: we thus observe that the SoBaP procedure exploiting the frequency dependence between the wavenumbers shows a faster decrease of the MSE as the SNR increases. For an SNR of 10 dB, the method reaches an MSE around -35 dB, compared to -20 dB for OMP.

This particular setting is considered in more detail in Fig. 5(b), where the performance at SNR = 10 dB is assessed from a sensor point of view. Within this setting, the three methods achieve similar (poor) performance for very small numbers of sensors (below $L = 50$). Note that this observation does not contradict the visual interpretation of Fig. 4, as the MSE quantifies the position of the $M[\nu]$ largest estimated wavenumbers without any consideration over false alarms. Beyond 50 sensors, the proposed SoBaP procedure outperforms indisputably the other two approaches.

B. North Sea Data

In this section, we apply our approach to real data. The considered data were acquired in the North Sea during a seismic campaign led by the Compagnie Générale de Géophysique [44], [45]. The source is an air gun that is impulsive with a nearly flat spectrum between 0 and 80 Hz. Measurements are performed by a synthetic antenna of 240 ocean-bottom seismometers resting on the seabed. They are spaced at intervals of 25 m, leading to a total length of 6000 m. The pressure field is sampled at 250 Hz. The environment is assumed to be close to a Pekeris waveguide. Within this assumption, the experimental parameters were estimated [45] as: $c_{\text{water}} = 1520$ m/s, $c_{\text{seabed}} = 1875$ m/s, and $D = 130$ m. For these data, we evaluated the SNR around 13 dB.

Fig. 6 presents the $f - k$ representations obtained by the inversion of 20 hydrophones with: a simple LS method [Fig. 6(b)], the OMP algorithm [Fig. 6(c)], and the SoBaP algorithm exploiting the frequency dependence between the wavenumbers [Fig. 6(d)]. As a comparison, the $f - k$ diagram obtained with

an LS inversion using the entire antenna is displayed in Fig. 6(a). Moreover, in each figure, we have added in red the wavenumbers selected from the LS inversion on the entire antenna [Fig. 6(a)] using a peak-picking method. These results show the relevance of the proposed approach. We can thus see that sparsity enables an undeniable improvement of the inversion (compared to the simple LS estimate) in particular in terms of noise suppression, but missed (in particular at the highest frequencies) and false detections are still noticeable. Adding prior knowledge on the propagated wavenumbers allows us to recover perfectly the wavenumbers, even at the highest frequencies (around 60 Hz in our case), while reducing the noise of the $f - k$ diagram.

V. CONCLUSION

This paper deals with the estimation of the wavenumbers in shallow-water environments. An array processing method is proposed to reduce the number of the sensors required to separate accurately the wavenumbers over a wide range of frequencies. The proposed method relies on two robust hypotheses. On the one hand, in the wavenumber domain, the wavenumber spectrum is sparse. On the other hand, in the frequency domain, the wavenumber spectra can be related from one frequency to the next using a general dispersion relationship. These two hypotheses call for the utilization of a Bayesian compressed-sensing (CS) algorithm. Indeed, the CS framework is suitable for our sensing context, while the Bayesian framework enables a natural implementation of the relation linking the wavenumbers from one frequency to the next.

The performance of the proposed method is assessed on simulations. Our algorithm makes possible an accurate wavenumber estimation using a small number of hydrophones, even in noisy contexts. It outperforms other state-of-the-art methods such as LS SFT (which does not benefit from any physical hypothesis) and OMP (which benefits only from the sparse hypothesis): estimated wavenumbers are more accurate, and the estimated $f - k$ diagrams are globally less noisy. Also, the method is validated on experimental data recorded in the North Sea. In this context, the proposed method performs the estimation of the wavenumbers with a 20-hydrophone HLA.

Because the proposed method is based on CS, it is naturally suitable for cases where the number of measurements (i.e., hydrophones) is small and, ideally, randomly distributed. However, it does not necessarily help in reducing the required range aperture (i.e., the HLA length). As a result, the method may not show its full strength on a small HLA with constant hydrophone spacing. However, interesting perspectives arise for analyzing (existing) synthetic aperture data, where the considered snapshots can be chosen at will along the source/receiver track. Also, the potential of CS should be taken into account when designing new at-sea experiment.

REFERENCES


- [1] T. C. Yang, "Effectiveness of mode filtering: A comparison of matched-field and matched-mode processing," *J. Acoust. Soc. Amer.*, vol. 87, 1990.
- [2] G. Wilson, R. Koch, and P. Vidmar, "Matched mode localization," *J. Acoust. Soc. Amer.*, vol. 84, pp. 310–320, 1988.
- [3] J. Buck, J. Preisig, and K. Wage, "A unified framework for mode filtering and the maximum a posteriori mode filter," *J. Acoust. Soc. Amer.*, vol. 103, p. 1813, 1998.
- [4] T. Neilsen and E. Westwood, "Extraction of acoustic normal mode depth functions using vertical line array data," *J. Acoust. Soc. Amer.*, vol. 111, pp. 748–756, 2002.
- [5] G. Potty, J. Miller, J. Lynch, and K. Smith, "Tomographic inversion for sediment parameters in shallow water," *J. Acoust. Soc. Amer.*, vol. 108, pp. 973–986, 2000.
- [6] J. Bonnel, B. Nicolas, J. Mars, and S. Walker, "Estimation of modal group velocities with a single receiver for geoacoustic inversion in shallow water," *J. Acoust. Soc. Amer.*, vol. 128, pp. 719–727, 2010.
- [7] J. Bonnel, G. Le Touzé, B. Nicolas, and J. Mars, "Physics-based time-frequency representations for underwater acoustics: Power class utilization with waveguide-invariant approximation," *IEEE Signal Process. Mag.*, vol. 30, no. 6, pp. 120–129, 2013.
- [8] C. Gervaise, B. Kinda, J. Bonnel, Y. Stéphan, and S. Vallez, "Passive geoacoustic inversion with a single hydrophone using broadband ship noise," *J. Acoust. Soc. Amer.*, vol. 131, no. 3, pp. 1999–2010, 2012.
- [9] Y. Le Gall and J. Bonnel, "Passive estimation of the waveguide invariant per pair of modes," *J. Acoust. Soc. Amer.*, vol. 134, no. 2, pp. EL230–EL236, 2013.
- [10] G. Frisk and J. Lynch, "Shallow water waveguide characterization using the Hankel transform," *J. Acoust. Soc. Amer.*, vol. 76, p. 205, 1984.
- [11] G. V. Frisk *et al.*, "Modal mapping experiment and geoacoustic inversion using sonobuoys," *IEEE J. Ocean. Eng.*, vol. 40, no. 3, pp. 607–620, 2015.
- [12] Ö. Yilmaz, *Seismic Data Analysis: Processing, Inversion, and Interpretation of Seismic Data*. no. 10, Tulsa, OK, USA: Society of Exploration Geophysicists, 2001, ch. 1.2.
- [13] B. Nicolas, J. Mars, and J. Lacoume, "Geoacoustical parameters estimation with impulsive and boat-noise sources," *IEEE J. Ocean. Eng.*, vol. 28, no. 3, pp. 494–501, 2003.
- [14] S. M. Kay, "Modern spectral estimation, theory and application," in *Signal Processing*. Upper Saddle River, NJ, USA: Prentice-Hall, 1988, ch. 6.
- [15] E. C. Shang, H. P. Wang, and Z. Y. Huang, "Waveguide characterization and source localization in shallow water waveguides using the prony method," *J. Acoust. Soc. Amer.*, vol. 83, no. 1, pp. 103–108, 1988.
- [16] K. M. Becker and G. V. Frisk, "Evaluation of an autoregressive spectral estimator for modal wave number estimation in range-dependent shallow water waveguides," *J. Acoust. Soc. Amer.*, vol. 120, no. 3, pp. 1423–1434, 2006.
- [17] F. D. Philippe, P. Roux, and D. Cassereau, "Iterative high-resolution wavenumber inversion applied to broadband acoustic data," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 55, no. 10, pp. 2306–2311, 2008.
- [18] F. Le Courtois and J. Bonnel, "Autoregressive model for high-resolution wavenumber estimation in a shallow water environment using a broadband source," *J. Acoust. Soc. Amer.*, vol. 135, no. 4, pp. EL199–EL205, 2014.
- [19] I.-T. Lu, R. C. Qiu, and J. Kwak, "A high-resolution algorithm for complex spectrum search," *J. Acoust. Soc. Amer.*, vol. 104, no. 1, pp. 288–299, 1998.
- [20] S. D. Rajan and S. D. Bhatta, "Evaluation of high-resolution frequency estimation methods for determining frequencies of eigenmodes in shallow water acoustic field," *J. Acoust. Soc. Amer.*, vol. 93, no. 1, pp. 378–389, 1993.
- [21] E. J. Candès and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 21–30, 2008.
- [22] N. R. Chapman and I. Barrodale, "Deconvolution of marine seismic data using the l_1 norm," *Geophys. J. Int.*, vol. 72, no. 1, pp. 93–100, 1983.
- [23] P. Gerstoft, A. Xenaki, and C. F. Mecklenbräuker, "Single and multiple snapshot compressive beamforming," arXiv:1503.02339, 2015.
- [24] J. B. Harley and J. M. Moura, "Dispersion curve recovery with orthogonal matching pursuit," *J. Acoust. Soc. Amer.*, vol. 137, no. 1, pp. EL1–EL7, 2015.
- [25] F. Le Courtois and J. Bonnel, "Compressed sensing for wideband wavenumber tracking in dispersive shallow water," *J. Acoust. Soc. Amer.*, vol. 138, no. 2, pp. 575–583, 2015.
- [26] F. Le Courtois and J. Bonnel, "Wavenumber tracking in a low resolution frequency-wavenumber representation using particle filtering," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, Florence, Italy, Jun. 2014, pp. 6805–6809.
- [27] P. M. Djuric *et al.*, "Particle filtering," *IEEE Signal Process. Mag.*, vol. 20, no. 5, pp. 19–38, Sep. 2003.
- [28] C. Soussen, J. Idier, D. Brie, and J. Duan, "From Bernoulli-Gaussian deconvolution to sparse signal restoration," *IEEE Trans. Signal Process.*, vol. 59, no. 10, pp. 4572–4584, 2011.

- [29] A. Drémeau, C. Herzet, and L. Daudet, "Boltzmann machine and mean-field approximation for structured sparse decompositions," *IEEE Trans. Signal Process.*, vol. 60, no. 7, pp. 3425–3438, 2012.
- [30] F. Jensen, W. Kuperman, M. Porter, and H. Schmidt, *Computational Ocean Acoustics*. New York, NY, USA: American Institute of Physics, 2nd ed., 2011, ch. 5 and 10.
- [31] E. J. Candès and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 21–30, 2008.
- [32] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [33] B. K. Natarajan, "Sparse approximate solutions to linear systems," *SIAM J. Comput.*, vol. 24, pp. 227–234, Apr. 1995.
- [34] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. Signal Process.*, vol. 41, pp. 3397–3415, Dec. 1993.
- [35] Y. C. Pati, R. Rezaifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Proc. 27th Annu. Asilomar Conf. Signals Syst. Comput.*, Nov. 1993, pp. 40–44.
- [36] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.*, vol. 20, pp. 33–61, 1998.
- [37] R. Tibshirani, "Regression shrinkage and selection via the LASSO," *J. Roy. Stat. Soc.*, vol. 58, pp. 267–288, 1996.
- [38] B. A. Olshausen and D. J. Field, "Sparse coding with an overcomplete basis set: A strategy employed by v1?" *Vis. Res.*, vol. 37, no. 23, pp. 3311–3325, 1997.
- [39] F. Krzakala, M. Mézard, F. Sausset, Y. F. Sun, and L. Zdeborová, "Statistical physics-based reconstruction in compressed sensing," 2011. [Online]. Available at: <http://arxiv.org/abs/1109.4424v2>
- [40] C. Herzet and A. Drémeau, "Bayesian pursuit algorithms," in *Proc. 18th Eur. Signal Process. Conf.*, Aug. 2010, pp. 1474–1478.
- [41] M. Beal, "Variational algorithms for approximate Bayesian inference," Ph.D. dissertation, Univ. College London, London, U.K., May 2003.
- [42] F. Krzakala, A. Manoel, E. W. Tramel, and L. Zdeborová, "Variational free energies for compressed sensing," in *Proc. IEEE Int. Symp. Inf. Theory*, Jul. 2014, pp. 1499–1503.
- [43] Y. C. Pati, R. Rezaifar, and P. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Conf. Rec. IEEE 27th Asilomar Conf. Signals Syst. Comput.*, 1993, pp. 40–44.
- [44] J. Bonnel, B. Nicolas, and J. I. Mars, "Estimation of modal group velocities with a single receiver for geoacoustic inversion in shallow water," *J. Acoust. Soc. Amer.*, vol. 128, no. 2, pp. 719–727, 2010.
- [45] B. Nicolas, J. I. Mars, and J.-L. Lacoume, "Geoacoustical parameters estimation with impulsive and boat-noise sources," *IEEE J. Ocean. Eng.*, vol. 28, no. 3, pp. 494–501, 2003.



Angélique Drémeau received the State Engineering degree from Telecom Bretagne, Brest, France, in 2007 and the M.Sc. and the Ph.D. degrees in signal processing and telecommunications from the Université de Rennes, Rennes, France, in 2007 and 2010, respectively.

She is currently an Associate Professor at ENSTA Bretagne, Lab-STICC (UMR 6285), Brest, France. Her research interests include inverse problems, sparse representation algorithms, and underwater acoustics.



Florent Le Courtois received the Ph.D. degree in acoustics from Université du Maine, Le Mans, France, in 2012.

From 2008 to 2011, he worked for the SNCF, Paris, France, on the problem of imaging acoustic sources on high-speed trains. From 2012 to 2015, he was employed as a Postdoctoral Fellow to work on the modal estimation using hydrophone arrays at ENSTA Bretagne, Brest, France. In January 2016, he joined Shom, Brest, France, to develop the underwater noise pollution descriptor for the MSFD.

His research interests include array processing, time-frequency analysis, sound-scapes, and bioacoustics.



Julien Bonnel (S'08–M'11) received the Ph.D. degree in signal processing from Grenoble Institut National Polytechnique (Grenoble INP), Grenoble, France, in 2010.

Since 2010, he has been an Associate Professor at Lab-STICC (UMR 6285), ENSTA Bretagne, Brest, France. His research in signal processing and underwater acoustics includes time-frequency analysis, source detection/localization, geoacoustic inversion, acoustical tomography, passive acoustic monitoring, and bioacoustics.

Dr. Bonnel is a member of the Acoustical Society of America.

Modal estimation in underwater acoustics by data-driven structured sparse decompositions

Clément Dorffer*, Thomas Paviet-Salomon*, Gilles Le Chenadec* and Angélique Drémeau*

*Lab-STICC & ENSTA-Bretagne, UMR CNRS 6285, Brest, France.

E-mail: clement.dorffer@ensta-bretagne.fr, thomas.paviet-salomon@ensta-bretagne.org,
gilles.le_chenadec@ensta-bretagne.fr, angelique.dremeau@ensta-bretagne.fr

Abstract—In underwater acoustics, shallow water environments act as modal dispersive waveguides when considering lowfrequency sources. In this context, propagating signals can be described as a sum of few modal components, each of them propagating according to its own wavenumber. Estimating these wavenumbers is of key interest to understand the propagating environment as well as the emitting source. To solve this problem, we proposed recently a Bayesian approach exploiting a sparsity-inforcing prior. When dealing with broadband sources, this model can be further improved by integrating the particular dependence linking the wavenumbers from one frequency to the other. In this contribution, we propose to resort to a new approach relying on a restricted Boltzmann machine, exploited as a generic structured sparsity-inforcing model. This model, derived from deep Bayesian networks, can indeed be efficiently learned on physically realistic simulated data using well-known and proven algorithms.

Index Terms—Underwater modal estimation, structured sparse approximation, Restricted Boltzmann Machine, Bayesian algorithm.

I. INTRODUCTION

In underwater acoustics, shallow environments behave like dispersive waveguides when considering low-frequency sources. An acoustic field received on an antenna is then classically described by a small set of modes propagating longitudinally according to their horizontal wavenumbers. The knowledge of these modes is of great importance for the characterization of the observation environment and, consequently, for the source localization. Among the different methods used to discriminate these modal components, the *frequency-wavenumber* (f-k) representation (see Fig. 1) allows a direct observation of the dispersion (*i.e.*, the frequency dependence) of the wavenumbers. Inherently conceivable for a horizontal array of sensors aligned with the source, they are particularly used in geophysics [1]. Recent contributions have focused on the construction of (f-k) diagrams by exploiting less constrained acquisition schemes, allowing their use in underwater acoustics.

Since propagation is described by a small number of modes, the use of sparsity-inforcing models seems appropriate. In fact, some contributions (see *e.g.*, [2], [3]) have proposed the use of the “compressed sensing” paradigm to estimate modal dispersion. However, if these methods prove to be relevant, we argue that they can be further improved by precisely integrating the dispersion relation linking the wavenumbers

Thanks to DGA Naval Systems for funding.




Fig. 1. Illustration of a (f-k) diagram obtained with a Pekeris waveguide.

from one frequency to another into the estimation process of the (f-k) diagram. To this end, we propose to build on some recent contributions in compressed sensing and machine learning.

On the one hand, research on compressed sensing [4]–[6] and more generally sparse decompositions have underlined the interest of taking into account the structures naturally living in signal representations [7]. On the other hand, restricted Boltzmann machines (RBM) [8] are at the heart of many recent contributions in machine learning that emphasize two precious qualities: *i)* they have been identified as generic models approximating any distribution over $\{0, 1\}^n$ [9]; *ii)* efficient algorithms have been developed to train them making them powerful representational models when large data sets are available [10]–[12].

These parallel researches in finally not so distant fields have recently led to the idea of exploiting RBMs learned from large databases to model the structures of sparse representations [13]–[15]. This paper follows on from these contributions. More particularly, we propose to deal with the (f-k) diagram estimation using a new Bayesian algorithm. Our approach considers a RBM as a model for structured sparse representations and exploits it through a mean-field approximation.

II. ACOUSTIC PROPAGATION IN SHALLOW WATER ENVIRONMENTS

In shallow water, acoustic propagation is described by modal theory. According to the latter, when considering a source at depth z_s and frequency f , the signal received by an hydrophone placed at depth z and distance r of the source can be formulated as

$$y(f, r, z, z_s) = \sum_{m=1}^{K(f)} A_m(f, z, z_s) e^{-ir k_{rm}(f)} + n(f, r, z), \quad (1)$$

where $K(f)$ stands for the number of modal components propagating at frequency f , $A_m(f, z, z_s)$ is the modal amplitude associated to the m -th component, $k_{rm}(f)$ is the m -th wavenumber at frequency f and $n(f, r, z)$ stands for some additive measurement noise.

For a M sensor linear antenna, Eq. (1) can be written as

$$\mathbf{y}_f = \mathbf{F}\mathbf{z}_f + \mathbf{n}_f \quad (2)$$

where \mathbf{F} is a $(M \times N)$ -dictionary of Fourier discrete atoms, \mathbf{z}_f contains a few non-zero coefficients indicating the wavenumbers propagating at frequency f and quantifying their amplitudes and \mathbf{n}_f stands for the additive noise.

A common approach to estimate \mathbf{z}_f from \mathbf{y}_f is to use a simple (inverse) spatial Fourier transform. Although very popular, the method requires sensors that are finely spaced sufficiently to avoid aliasing. As an alternative, more recent contributions [2], [16], [17] have proposed to leverage the sparsity of \mathbf{z}_f to design more robust methods.

When dealing with broadband sources, the generative model can be written as follows :

$$[\mathbf{y}_{f_1}, \dots, \mathbf{y}_{f_P}] = \mathbf{F} \cdot [\mathbf{z}_{f_1}, \dots, \mathbf{z}_{f_P}] + [\mathbf{n}_{f_1}, \dots, \mathbf{n}_{f_P}], \quad (3)$$

where P is the number of (discretized) frequencies of the source. The estimation of the \mathbf{z}_f 's can then be performed frequency per frequency with the same tools as for the mono-frequency case. Stacking these estimates \mathbf{z}_f one on top of the other leads to the (f-k) diagram (see Fig. 1) and makes appear some structures linking the wavenumbers propagating from one frequency to the other. These structures are well-known in practice under the physical name of dispersion relation. While some contributions in the literature work on an explicit combination of this relation with sparse approaches [2], [17], we propose here to resort to a Bayesian formulation of the (f-k) diagram estimation problem using a particular choice of prior holding the structure information.

To that end, a vectorized version of Eq. (3) is built by defining $\mathbf{y} = [\mathbf{y}_{f_1}^T, \dots, \mathbf{y}_{f_P}^T]^T$, $\mathbf{z} = [\mathbf{z}_{f_1}^T, \dots, \mathbf{z}_{f_P}^T]^T$, $\mathbf{n} = [\mathbf{n}_{f_1}^T, \dots, \mathbf{n}_{f_P}^T]^T$ and \mathbf{D} as a block-diagonal matrix with \mathbf{F} P -times repeated on its diagonal. Adopting this formulation, Eq. 3 can be re-formulated as

$$\mathbf{y} = \mathbf{D}\mathbf{z} + \mathbf{n}, \quad (4)$$

where \mathbf{y} (resp. \mathbf{z}) is of dimension MP (resp. NP) and \mathbf{D} is a $(MP \times NP)$ -dictionary.

The inverse problem we are interested in is then formulated as the estimation of the (f-k) diagram \mathbf{z} from the acoustic signal \mathbf{y} measured over the M -sensor antenna at all of the P frequencies.

III. PROBLEM FORMULATION

According to the modal theory, in shallow water environments and at low frequencies, the vector \mathbf{z} has few non-zero elements, corresponding to the propagating modal wavenumbers, that is \mathbf{z} is assumed to be sparse. Modeling such a property in a Bayesian framework can be realized in different ways. A popular model considers \mathbf{z} as the result of the Hadamard product of a centered circular Gaussian variable¹, say $\mathbf{x} \in \mathbb{C}^{NP}$ with variance σ_x^2 and a binary variable, say $\mathbf{s} \in \{0, 1\}^{NP}$. This model explicitly expresses the *support* \mathbf{s} of the sparse vector \mathbf{z} , which is a key component of sparse estimation²: we look for a good estimation of \mathbf{s} knowing the observations \mathbf{y} .

When no structure is assumed in sparse representations, the Bernoulli law constitutes a natural choice for \mathbf{s} . Structures, for their part, can be modeled according to various distributions. In [7], [18]–[20], Boltzmann machines are envisaged as generic models encompassing many well-known models. These models depend on parameters which are difficult to train, due - among others - to the presence of a quadratic term. Instead, more recent works [13]–[15] propose to exploit restricted Boltzmann machines, for which efficient training algorithms exist.

Formally, for $\mathbf{s} \in \{0, 1\}^{NP}$, RBM can be expressed as

$$p(\mathbf{s}) = \sum_{\mathbf{h}} p(\mathbf{s}, \mathbf{h}) \propto \exp(\mathbf{a}^T \mathbf{h} + \mathbf{b}^T \mathbf{s} + \mathbf{s}^T \mathbf{W} \mathbf{h}), \quad (5)$$

where \mathbf{h} is a L -dimensional binary hidden variable and \mathbf{a} , \mathbf{b} and \mathbf{W} are the RBM parameters.

In [14], the authors propose to implement this model into a OMP-like framework, while in [15], a reweighted ℓ_1 -like procedure is considered. Both procedures are deterministic, integrating RBM as an additional extra block and making hard decision on it. In [13], a Bayesian approach based on an approximate message passing procedure is proposed, leading to a more integrated vision of the restricted Boltzmann machine. However, RBM is not fully exploited in the sense that the support \mathbf{s} is still assumed to follow a Bernoulli law, although informed by the RBM. In this paper, we propose a new Bayesian approach, leveraging both *i*) the probabilistic nature of the RBM, and *ii*) the explicit modeling of \mathbf{s} as the commonly visible layer of a RBM (see Eq. (5)).

More particularly, considering model (4)-(5), we are interested in the following marginalized maximum a posteriori estimation problem :

$$(\hat{\mathbf{x}}, \hat{\mathbf{s}}) = \underset{\mathbf{x}, \mathbf{s}}{\operatorname{argmax}} \log p(\mathbf{x}, \mathbf{s} | \mathbf{y}) \quad (6)$$

¹For a sake of simplicity, we will use the same notation for a random variable and its realizations.

²We note indeed that once the support is estimated, the corresponding coefficients emerge straightforwardly by least-squares estimation.

with $p(\mathbf{x}, \mathbf{s}|\mathbf{y}) = \int p(\mathbf{x}, \mathbf{s}, \mathbf{h}|\mathbf{y})$. In the continuation of [7], we propose to resort to a mean-field approximation.

IV. DEEP STRUCTURED SOFT BAYESIAN PURSUIT

Mean-field (MF) approximations aim at approximating a posterior distribution, here $p(\mathbf{x}, \mathbf{s}|\mathbf{y})$, by a ‘‘simpler’’ distribution, say $q(\mathbf{x}, \mathbf{s})$, having a easy-to-handle factorization. We consider in this paper the factorization

$$q(\mathbf{x}, \mathbf{s}, \mathbf{h}) = \prod_{i=1}^{NP} q(x_i, s_i) \prod_{j=1}^L q(h_j), \quad (7)$$

with x_i (resp. s_i) the i -th element in \mathbf{x} (resp. \mathbf{s}) and h_j the j -th element in \mathbf{h} . The approximation $q(\mathbf{x}, \mathbf{s}, \mathbf{h})$ is chosen to be *as close as possible* to $p(\mathbf{x}, \mathbf{s}, \mathbf{h}|\mathbf{y})$ in the sense of the Kullback-Leibler (KL) divergence. In practice, this latter optimization problem can be efficiently solved by an iterative algorithm, called ‘‘variational Bayesian expectation maximization’’ (VBEM) algorithm, insuring a decreasing of the KL divergence at each iteration. We expose here below its particularization to model (4)-(5).

The VBEM procedure successively updates all factors in the MF approximation, namely here the $q(x_i, s_i)$'s and $q(h_j)$'s, according to the following iterative rules:

$$q^{(k+1)}(x_i, s_i) \propto \exp \left(\langle \log p(\mathbf{x}, \mathbf{s}, \mathbf{h}, \mathbf{y}) \rangle_{\substack{\prod_j q^{(k)}(h_j) \\ \prod_{j>i} q^{(k)}(x_j, s_j) \\ \prod_{j<i} q^{(k+1)}(x_j, s_j)}} \right)$$

$$q^{(k+1)}(h_i) \propto \exp \left(\langle \log p(\mathbf{x}, \mathbf{s}, \mathbf{h}, \mathbf{y}) \rangle_{\substack{\prod_{j>i} q^{(k)}(h_j) \\ \prod_{j<i} q^{(k+1)}(h_j) \\ \prod_j q^{(k+1)}(x_j, s_j)}} \right)$$

where $\langle f(\mathbf{u}) \rangle_{q(u_j)} \triangleq \int_{u_j} q(u_j) f(\mathbf{u}) du_j$ and k is the current iteration number. Developing these update rules according to model (4)-(5) gives $q^{(k+1)}(x_i, s_i) = q^{(k+1)}(x_i|s_i)q^{(k+1)}(s_i)$

$$\text{with } q^{(k+1)}(x_i|s_i) = \mathcal{CN}(m_x^{(k+1)}(s_i), \Sigma_x^{(k+1)}(s_i)) \quad (8)$$

$$m_x^{(k+1)}(s_i) = s_i \frac{\sigma_x^2}{\sigma_n^2 + \sigma_x^2 \mathbf{d}_i^H \mathbf{d}_i} \mathbf{d}_i^H \langle \mathbf{r}_i \rangle^{(k+1)} \quad (9)$$

$$\Sigma_x^{(k+1)}(s_i) = \frac{\sigma_n^2 \sigma_x^2}{\sigma_n^2 + \sigma_x^2 \mathbf{d}_i^H \mathbf{d}_i} \quad (10)$$

$$\langle \mathbf{r}_i \rangle^{(k+1)} = \mathbf{y} - \sum_{j \neq i} q^{(n)}(s_j=1) m_x^{(n)}(s_j=1) \mathbf{d}_j \quad (11)$$

with $n=k+1$ if $j < i$ and $n=k$ if $j > i$

where $\mathcal{CN}(\mu, \Gamma)$ stands for the circular Gaussian distribution with mean μ and variance Γ and

$$q^{(k+1)}(s_i) \propto \exp \left(s_i \left(b_i + \sum_l w_{il} q^{(k)}(h_l = 1) \right) \right)$$

$$\sqrt{\Sigma_x^{(k+1)}(s_i)} \exp \left(\frac{1}{2} \frac{|m_x^{(k+1)}(s_i)|^2}{\Sigma_x^{(k+1)}(s_i)} \right) \quad (12)$$

$$q^{(k+1)}(h_l) \propto \exp \left(h_l \left(a_l + \sum_i w_{il} q^{(k+1)}(s_i = 1) \right) \right) \quad (13)$$

where we have all along denoted by \mathbf{d}_i the i -th column in \mathbf{D} and w_{il} the (i, l) -th element in \mathbf{W} .

To help the convergence of the algorithm, an estimation of the noise variance σ_n^2 is implemented in the same way as the one proposed in SoBaP [7], leading to a dependence of $\Sigma_x^{(k+1)}(s_i)$ in the iteration number $(k+1)$ (see Eq. (10)). As stopping criterion, the KL divergence is computed at each iteration, between the target distribution $p(\mathbf{x}, \mathbf{s}, \mathbf{h}|\mathbf{y})$ and its current MF approximate $q(\mathbf{x}, \mathbf{s}, \mathbf{h})$: the algorithm stops when this divergence no longer decreases ‘‘sufficiently’’. Pseudo-code 1 gives a practical implementation of the above equations. The use of RBMs being a natural bridge towards deep networks, we will refer to the proposed procedure as the ‘‘Deep Structured Soft Bayesian Pursuit’’ (DSSoBaP).

Pseudo-code 1 DSSoBaP algorithm

Input: $\mathbf{y}, \mathbf{D}, \mathbf{W}, \mathbf{a}, \mathbf{b}, \sigma_x^2, \sigma_n^2$
Initialisation: $\{q^{(0)}(s_i), m_x^{(0)}(s_i)\}_{s_i \in \{0,1\}, i \in \{1, \dots, NP\}}$,
 $\{q^{(0)}(h_l)\}_{l \in \{1, \dots, L\}}$, $k = 0$, $\text{KL} = 0$, $\text{KL}_{\text{old}} = \infty$
1: **while** $\text{KL}_{\text{old}} - \text{KL} > 10^{-1}$ **do**
2: **Optional:** update σ_n^2 according to [7]
3: **for** $i = 1 \dots NP$ **do**
4: update $\langle \mathbf{r}_i \rangle^{(k+1)}$ using (11)
5: **Optional:** if step 2, update $\Sigma_x^{(k+1)}(s_i)$ using (10)
6: update $m_x^{(k+1)}(s_i)$ using (9) for $s_i \in \{0, 1\}$
7: update $q^{(k+1)}(s_i)$ using (12) for $s_i \in \{0, 1\}$
8: **end for**
9: **for** $l = 1 \dots L$ **do**
10: update $q^{(k+1)}(h_l)$ using (13) for $h_l \in \{0, 1\}$
11: **end for**
12: $k \leftarrow k + 1$
13: $\text{KL}_{\text{old}} \leftarrow \text{KL}$
14: compute current KL divergence
15: **if** $\text{KL}_{\text{old}} - \text{KL} \leq 10^{-1}$ **then**, $k_{\text{max}} = k$
16: **end if**
17: **end while**
Output: $\{q^{(k_{\text{max}})}(s_i = 1), m_x^{(k_{\text{max}})}(s_i = 1)\}_{i \in \{1, \dots, NP\}}$

Note that DSSoBaP has an algorithmic complexity in the order of $\mathcal{O}(NP(MP + L))$ per iteration. This complexity is reasonable compared to other procedures in the literature, such as *e.g.* RBM-OMP [14]. This drawback in particular prevented us from using RBM-OMP in our particular problem.

V. EXPERIMENT

In this section, we confront our approach to synthetic experiments and compare its performance to the SoBaP standard procedure.

A. Data simulation and training phase

Several models of the ocean have been proposed in the literature of modal theory [21]. One of the most popular is the Pekeris waveguide. In this model, the sea surface is assumed to be perfectly reflective, and the sea bottom is considered as a semi-infinite fluid. Its parameters therefore




Fig. 2. (f-k) diagrams obtained for $M/N = 0.1$ and SNR= 3 dB by (left) inverse Fourier transform; (middle) SoBaP; (right) DSSoBaP.

include water layer thickness D (constant over the source-receiver distance), water density ρ_w and celerity c_w , ground density ρ_g and celerity c_g . Considering this simple model, we propose to simulate the propagation of an acoustic signal in various Pekeris environments and to learn the structures in the corresponding (f-k) diagrams. The considered setup is as follows.

Source and antenna are assumed to lie on the ground at same depth $z_s = z = D$ varying from 100 m to 200 m. We simulate the water celerity c_w from 1450 m s^{-1} to 1550 m s^{-1} , and the water density ρ_w from 950 kg m^{-3} to 1050 kg m^{-3} . The studied frequencies range from 0 Hz to 30 Hz with a step-size of 0.5 Hz, *i.e.*, $P = 60$ frequencies are considered, and we set the ground properties c_g and ρ_g from the sediments geoacoustic parameters synthesis in [22]. The simulated antenna is composed of $M = 120$ sensors regularly spaced from 50 m.

With such parameters, we generated a set of 3575 signals – of size 120×60 , *i.e.*, 120 sensors, 60 frequencies–and their theoretical sparse decomposition (f-k) diagrams – of size 120×60 , *i.e.*, $N = 120$ potential wavenumbers for each 60 frequencies – whose sparsity ratios fluctuate between $3 \cdot 10^{-3}$ and $3 \cdot 10^{-2}$. We then randomly divided these simulations into a training set (3000 samples) and a testing set (575 samples) and binarized the (f-k) diagrams from the training set so as to obtain their theoretical supports.

These supports were used to train a RBM with $L = 30$ hidden units. Training parameters were fixed to 10^{-2} for the learning rate, 50 for the minibatch size, 2000 for the number of epochs and the Persistent Contrastive Divergence algorithm [10] was employed with an l_1 regularization so as to enforce the sparsity of the learned features. The obtained RBM is then integrated as prior into DSSoBaP.

B. Tests and results analysis

To assess the relevance of RBM into the considered inverse problem, we compare the performance of DSSoBaP with that of SoBaP, introduced in [7]. This algorithm is based on the same MF approximation as DSSoBaP but instead of the RBM prior (see Eq. (5)), SoBaP exploits a *i.i.d.* Bernoulli model on the SR support. It has to be noticed that we do not compare

the approach with those proposed in [7], [14], mainly because of their computational cost.

Two figures of merit are considered: *i*) the estimation of the modal amplitudes (*i.e.*, the non-zero coefficients in the \mathbf{z}_f 's) is evaluated according to the normalized mean square error (NMSE), *ii*) the detection of the wavenumbers (*i.e.*, the support of the sparse \mathbf{z}_f 's) is evaluated according to the True Positives Rate (TPR) which is the proportion of the theoretical support that has been correctly recovered and the False Discovery Rate (FDR) which is the proportion of the estimated support that should not have been activated (*i.e.*, the proportion of false detection within the estimated support). Each metric is valued according to the signal-to-noise ratio (SNR) on the one hand, and the number of sensors on the other hand. For the tests related to the SNR, we fix the number of observations to 1440 by randomly selecting 24 sensors and let the SNR range from 0 dB to 5 dB by adding Gaussian noise \mathbf{n}_f with modulated variance σ_n^2 . For the tests related to the number of sensors – which is classical in compressed sensing – we choose to fix the SNR to 3 dB and let the proportion of hydrophones range from 0.1 to 1 by randomly selecting from 12 to 120 sensors of the antenna.

SoBaP and DSSoBaP are run using the same stopping criterium as described in pseudo-code 1. For SoBaP, the Bernoulli parameter is set for each test sample to its exact value that is the theoretical proportion of nonzero coefficients in \mathbf{z} , while for DSSoBaP the prior is provided by the same trained RBM for all test samples. Finally, the performance metrics are averaged over the entire test dataset.

Fig. 3 illustrates the obtained results. One can see that DSSoBaP performance curves follow the same tendencies as SoBaP while being better in most cases, especially in the worst configurations, *i.e.*, with low SNR and low M/N ratio. In particular, the gap between DSSoBaP and SoBaP TPR curves shows that the RBM-based prior has been more effective than the Bernoulli prior for recovering the theoretical support, while the similar FDR curves show that it also conserves the same sparsity properties of the Bernoulli prior by producing only slightly more false detections. Interestingly, we can observe that the improvement of the detection performance leads to an improvement of the NMSE metric, that is the amplitudes of the modes are consequently better recovered with DSSoBaP




Fig. 3. NMSE (top) and TPR/FDR (bottom) w.r.t. the SNR (left) and the sensor ratio M/N (right).

than with SoBaP.

An example of recovered (f-k) diagrams by inverse Fourier transform, SoBaP and DSSoBaP is shown on Fig. 2. The considered test parameters are SNR=3 dB and sensor ratio $M/N = 0.1$. One can see that DSSoBaP is able to detect much elements from the ground truth support than SoBaP. In particular, structures that are characteristic of modal propagation – the line shapes induced by wavenumber propagation along frequencies – are better recognized using the RBM-based prior, thus showing that RBM are, as expected, well suited for learning spatial dependencies in the support density.

VI. CONCLUSION

In this work, we introduced a new Bayesian sparse decomposition algorithm able to exploit an RBM – learned in a previous step – as a structure model on the support of a sparse representation. The proposed algorithm relies on an MF approximation and allows a natural integration of the RBM. Its performance on a current underwater acoustic problem is very promising compared to its “standard” counterpart which exploits non-structured Bernoulli variables. Moreover, as the training phase is performed on a collection of physically-realistic simulations, it naturally incites the prior to respect some theoretical structures, which might be helpful when dealing with real datasets.

REFERENCES

- [1] L. Amundsen and B. Ursin, “Frequency-wavenumber inversion of acoustic data,” *Geophysics*, vol. 56, no. 7, pp. 1027–1039, 1991.
- [2] A. Drémeau, F. Le Courtois, and J. Bonnel, “Reconstruction of dispersion curves in the frequency-wavenumber domain using compressed sensing on a random array,” *IEEE Journal of Oceanic Engineering*, vol. 42, no. 4, pp. 914–922, 2017.
- [3] J. B. Harley and J. M. Moura, “Dispersion curve recovery with orthogonal matching pursuit,” *The Journal of the Acoustical Society of America*, vol. 137, no. 1, pp. EL1–EL7, 2015.

- [4] D. L. Donoho, “Compressed sensing,” *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [5] E. J. Candès and M. B. Wakin, “An introduction to compressive sampling,” *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21–30, 2008.
- [6] Y. C. Eldar and G. Kutyniok, *Compressed sensing: theory and applications*. Cambridge university press, 2012.
- [7] A. Drémeau, C. Herzet, and L. Daudet, “Boltzmann machine and mean-field approximation for structured sparse decompositions,” *IEEE Transactions On Signal Processing*, vol. 60, no. 7, p. 3425–3438, 2012.
- [8] P. Smolensky, “Information processing in dynamical systems: Foundations of harmony theory,” Colorado Univ at Boulder Dept of Computer Science, Tech. Rep., 1986.
- [9] N. Le Roux and Y. Bengio, “Representational power of restricted boltzmann machines and deep belief networks,” *Neural computation*, vol. 20, no. 6, pp. 1631–1649, 2008.
- [10] T. Tieleman, “Training restricted boltzmann machines using approximations to the likelihood gradient,” in *Proc. 25th international conference on Machine learning*, 2008, pp. 1064–1071.
- [11] G. E. Hinton, “A practical guide to training restricted boltzmann machines,” in *Neural networks: Tricks of the trade*. Springer, 2012, pp. 599–619.
- [12] A. Fischer and C. Igel, “Training restricted boltzmann machines: An introduction,” *Pattern Recognition*, vol. 47, no. 1, pp. 25–39, 2014.
- [13] E. W. Tramel, A. Manoel, F. Caltagirone, M. Gabrié, and F. Krzakala, “Inferring sparsity: Compressed sensing using generalized restricted boltzmann machines,” in *2016 IEEE Information Theory Workshop (ITW)*. IEEE, 2016, pp. 265–269.
- [14] L. F. Polania and K. E. Barner, “Exploiting restricted boltzmann machines and deep belief networks in compressed sensing,” *IEEE Transactions on Signal Processing*, vol. 65, no. 17, pp. 4538–4550, 2017.
- [15] Z. Liao, J. Zhang, D. Hu, C. Li, L. Zhu, and Y. Li, “Sparse signal reconstruction with statistical prior information: A data-driven method,” *IEEE Access*, vol. 7, pp. 157 037–157 045, 2019.
- [16] F. Le Courtois and J. Bonnel, “Compressed sensing for wideband wavenumber tracking in dispersive shallow water,” *The Journal of the Acoustical Society of America*, vol. 138, no. 2, pp. 575–583, 2015.
- [17] T. Paviet-Salomon, C. Dorffer, J. Bonnel, B. Nicolas, T. Chonavel, and A. Drémeau, “Dispersive grid-free orthogonal matching pursuit for modal estimation in ocean acoustics,” *Proc. Int’l Conference on acoustics, speech and signal processing*, 2019.
- [18] T. Peleg, Y. Eldar, and M. Elad, “Exploiting statistical dependencies in sparse representations for signal recovery,” *IEEE Transactions on Signal Processing*, vol. 60, no. 5, pp. 2286–2303, 2012.
- [19] V. Cevher, M. Duarte, C. Hegde, and R. Baraniuk, “Sparse signal recovery using markov random fields,” *Proc. Neural Information Processing Systems (NeurIPS)*, 2008.
- [20] J. Garrigues and B. A. Olshausen, “Learning horizontal connections in a sparse coding model of natural images,” *Proc. Neural Information Processing Systems (NeurIPS)*, 2008.
- [21] F. B. Jensen, W. A. Kuperman, M. B. Porter, and H. Schmidt, *Computational ocean acoustics*. Springer Science & Business Media, 2011.
- [22] D. Jackson, “High-frequency ocean environmental acoustic models handbook,” *Tech. report, Applied Physics Laboratory, Univ. of Washington*, 1994.

Reference-less measurement of the transmission matrix of a highly scattering material using a DMD and phase retrieval techniques

Angélique Drémeau,¹ Antoine Liutkus,² David Martina,³
Ori Katz,^{3,4} Christophe Schülke,⁵ Florent Krzakala,^{1,6}
Sylvain Gigan,^{3,6,*} and Laurent Daudet^{4,5}

¹LPS-ENS and CNRS UMR 8550, Paris, F-75005, France.

²Inria, CNRS, Loria UMR 7503, Villers-lès-Nancy, F-54600, France

³Laboratoire Kastler Brossel, Université Pierre et Marie Curie, Ecole Normale Supérieure,
Collège de France, CNRS UMR 8552, Paris, F-75005, France

⁴Institut Langevin, ESPCI and CNRS UMR 7587, Paris, F-75005, France

⁵Paris Diderot University, Sorbonne Paris Cité, Paris, F-75013, France

⁶Sorbonne Universités, UPMC Université Paris 06, F-75005, Paris, France

*sylvain.gigan@lkb.ens.fr

Abstract: This paper investigates experimental means of measuring the transmission matrix (TM) of a highly scattering medium, with the simplest optical setup. Spatial light modulation is performed by a digital micromirror device (DMD), allowing high rates and high pixel counts but only binary amplitude modulation. On the sensor side, without a reference beam, the CCD camera provides only intensity measurements. Within this framework, this paper shows that the TM can still be retrieved, through signal processing techniques of phase retrieval. This is experimentally validated on three criteria : quality of prediction, distribution of singular values, and quality of focusing.

© 2015 Optical Society of America

OCIS codes: (290.4210) Multiple scattering; (070.6120) Spatial light modulators; (100.5070) Phase retrieval

References and links

1. P. Sebbah, *Waves and Imaging Through Complex Media* (Springer, 2001).
2. A. P. Mosk, A. Lagendijk, G. Lerosey, and M. Fink, "Controlling waves in space and time for imaging and focusing in complex media," *Nature Photonics* **6**, 283–292 (2012).
3. M. Cui and C. Yang, "Implementation of a digital optical phase conjugation system and its application to study the robustness of turbidity suppression by phase conjugation," *Opt. Express* **18**, 3444–3455 (2010).
4. I. N. Papadopoulos, S. Farahi, C. Moser, and D. Psaltis, "Focusing and scanning light through a multimode optical fiber using digital phase conjugation," *Opt. Express* **20**, 10583 (2012).
5. I. M. Vellekoop and A. P. Mosk, "Focusing coherent light through opaque strongly scattering media," *Opt. Express* **32**, 2309 (2007).
6. S. M. Popoff, G. Lerosey, R. Carminati, M. Fink, A. C. Boccara, and S. Gigan, "Measuring the transmission matrix in optics: An approach to the study and control of light propagation in disordered media," *Phys. Rev. Lett.* **104**, 100601 (2010).
7. S. Popoff, G. Lerosey, M. Fink, A. C. Boccara, and S. Gigan, "Image transmission through an opaque material," *Nature Commun.* **1**, 81 (2010).
8. Y. Choi, T. D. Yang, C. Fang-Yen, P. Kang, K. J. Lee, R. R. Dasari, M. S. Feld, and W. Choi, "Overcoming the diffraction limit using multiple light scattering in a highly disordered medium," *Phys. Rev. Lett.* **107**, 023902 (2011).

9. M. Kim, Y. Choi, C. Yoon, W. Choi, J. Kim, Q.-H. Park, and W. Choi, "Maximal energy transport through disordered media with the implementation of transmission eigenchannels," *Nature Photonics* **6**, 581–585 (2012).
10. Y. Choi, C. Yoon, M. Kim, T. D. Yang, C. Fang-Yen, R. R. Dasari, K. J. Lee, and W. Choi, "Scanner-free and wide-field endoscopic imaging by using a single multimode optical fiber," *Phys. Rev. Lett.* **109** (2012).
11. I. N. Papadopoulos, S. Farahi, C. Moser, and D. Psaltis, "High-resolution, lensless endoscope based on digital scanning through a multimode optical fiber," *Biomed. Opt. Express* **4**, 260–270 (2013).
12. S. Bianchi and R. Di Leonardo, "A multi-mode fiber probe for holographic micromanipulation and microscopy," *Lab on a Chip* **12**, 635 (2012).
13. T. Čížmár and K. Dholakia, "Shaping the light transmission through a multimode optical fibre: complex transformation analysis and applications in biophotonics," *Opt. Express* **19**, 18871 (2011).
14. J. B. Sampson, "Dmd display system," (1995). US Patent 5,452,024.
15. D. B. Conkey, A. M. Caravaca-Aguirre, and R. Piestun, "High-speed scattering medium characterization with application to focusing light through turbid media," *Opt. Express* **20**, 1733–1740 (2012).
16. S. A. Goorden, J. Bertolotti, and A. P. Mosk, "Superpixel-based spatial amplitude and phase modulation using a digital micromirror device," arXiv:1405.3893 [physics] (2014).
17. D. Akbulut, T. J. Huisman, E. G. van Putten, W. L. Vos, and A. P. Mosk, "Focusing light through random photonic media by binary amplitude modulation," *Opt. Express* **19**, 4017–4029 (2011).
18. D. Kim, W. Choi, M. Kim, J. Moon, K. Seo, S. Ju, and W. Choi, "Implementing transmission eigenchannels of disordered media by a binary-control digital micromirror device," *Opt. Commun.* **330**, 35–39 (2014).
19. J. W. Tay, J. Liang, and L. V. Wang, "Amplitude-masked photoacoustic wavefront shaping and application in flowmetry," *Opt. Lett.* **39**, 5499–5502 (2014).
20. X. Zhang and P. Kner, "Binary wavefront optimization using a genetic algorithm," *Journal of Optics* **16**, 125704 (2014).
21. T. Chaigne, O. Katz, A. C. Boccara, M. Fink, E. Bossy, and S. Gigan, "Controlling light in scattering media non-invasively using the photoacoustic transmission matrix," *Nature Photonics* **8**, 58–64 (2014).
22. D. Akbulut, T. Strudley, J. Bertolotti, T. Zehender, E. Bakkers, A. Lagendijk, W. Vos, O. Muskens, and A. Mosk, "Measurements on the optical transmission matrices of strongly scattering nanowire layers," in *Proc. Conf. on Lasers and Electro-Optics Europe and International Quantum Electronics Conf. (CLEO EUROPE/IQEC)*, (2013).
23. A. Drémeau and F. Krzakala, "Phase recovery from a bayesian point of view: the variational approach," in *Proceedings of IEEE Trans. Acoust. Speech Signal Process.* (2015).
24. R. Gerchberg and W. Saxton, "A practical algorithm for the determination of phase from image and diffraction plane pictures," *Optik* **35**, 237–246 (1972).
25. J. R. Fienup, "Phase retrieval algorithms: a comparison," *Appl. Opt.* **21**, 2758–2769 (1982).
26. E. Candès, T. Strohmer, and V. Voroninski, "Phaselift : exact and stable signal recovery from magnitude measurements via convex programming," *Commun. in Pure and Applied Mathematics* **66**, 1241–1274 (2013).
27. I. Waldspurger, A. d'Aspremont, and S. Mallat, "Phase recovery, maxcut and complex semidefinite programming," *Mathematical Programming Series A- Springer* (2013).
28. P. Schniter and S. Rangan, "Compressive phase retrieval via generalized approximate message passing," in *Proceedings of Communication, Control, and Computing (Allerton)* (2012).
29. V. A. Marčenko and L. A. Pastur, "Distribution of eigenvalues for some sets of random matrices." *Mathematics of the USSR-Sbornik* **1**, 457 – 483 (1967).
30. S. M. Popoff, G. Lerosey, M. Fink, A. C. Boccara, and S. Gigan, "Controlling light through optical disordered media: transmission matrix approach," *New Journal of Physics* **13**, 123021 (2011).
31. A. Liutkus, D. Martina, S. Popoff, G. Chardon, O. Katz, G. Lerosey, S. Gigan, L. Daudet, and I. Carron, "Imaging with nature: Compressive imaging using a multiply scattering medium," *Sci. Rep.* **4** (2014).

1. Introduction

Wave propagation in complex media is a fundamental problem in physics, be it in acoustics, optics, or electromagnetism [1]. In optics, it is particularly relevant for imaging applications. Indeed, when light passes through a multiply scattering medium, such as a biological tissue or a layer of paint, ballistic light is rapidly attenuated, preventing conventional imaging techniques, and random scattering events generate a so-called speckle pattern that is usually considered useless for imaging. Recently, wavefront shaping using spatial light modulators (SLM) has emerged as a unique tool to manipulate multiply scattered coherent light, for focusing or imaging in scattering media [2]. In essence, these methods use the linearity and time-reversal symmetry of the wave propagation, whatever the complexity of the medium, to control the output speckle field, by manipulating the light beam impinging on the scattering sample. Dif-

ferent wavefront shaping approaches rely on digital phase-conjugation [3, 4] or iterative algorithms [5], but it is also possible to measure the so-called transmission matrix (TM) of the medium [6], which fully describes light propagation through the linear medium, from the modulator device to the detector. This approach has been particularly efficient for focusing, imaging [7, 8] and for studying the transmission modes of the medium [9]. These methods are not only valid for scattering material but can also be applied to other complex transmission system, most notably multimode fibers, turning them into minimal footprint endoscopes [10–13].

A major limitation of most of these techniques for imaging is their speed. Indeed, the wavefront shaping process must be faster than the stability time of the medium, which can be of only a few milliseconds in biological tissues. Yet, most of the works reported so far have relied on phase modulators which are usually slow (few tens of Hertz for liquid crystal modulators). Micro Electro-Mechanical Systems (MEMS) modulators are much faster, but are usually not phase-only. As a promising alternative for wave shaping in complex media, Digital Micromirror Device (DMD) technology [14] offers binary amplitude modulators (*i.e.*, “ON” or “OFF”) operating at $> 20kHz$, with high pixel counts (10^6) and low pitch (around 10 microns), all this at low cost. These binary amplitude modulators have been used as phase modulators, using appropriate diffraction and filtering, *e.g.* by Lee-type amplitude holography [15, 16], as shown on Fig. 1(b). While phase control is more effective for wavefront shaping than amplitude control, some works reported on using DMD as genuine binary amplitude modulators for wavefront shaping through opaque scattering media, albeit usually yielding lower overall efficiency than phase modulators for focusing or mode matching [17–19]. The DMD configuration can also be optimized using genetic algorithms [20] to maximize the intensity enhancement.

For the measurement of a TM, an additional issue lies in accessing the amplitude and phase of the output field, that in optics usually requires a holographic measurement, *i.e.*, a reference beam, as shown on Fig. 1(a). This reference beam can either be co-propagating in the medium [7, 21], or use an external reference arm [8, 22]. The phase and amplitude of the measured field can then be extracted by simple linear combinations of interference patterns with a phase-shifted or off-axis reference. This however poses the unavoidable experimental problem of the interferometric stability of the reference arm.

In this work we report on the full measurement of the complex TM of a multiply scattering medium, using a DMD binary amplitude modulator as an SLM, with no reference on the detection side, as shown on Fig. 1(c). This approach combines the high-speed and high pixel counts allowed by DMD devices, with the simplicity and robustness of a reference-less optical setup. However, it involves advanced signal processing algorithms for phase retrieval, run on a sufficiently large number of input-output calibration measurements. Here, we use a Bayesian phase retrieval algorithm [23] for the estimation of a TM, based on actual noisy experimental measurements. This is done one output pixel at a time, a process that can easily be parallelized for speed. Although a number of phase estimation techniques could also be used, our approach provides good results at moderate computational costs. This technique is validated on experimental data, with three criteria. First, for a given arbitrary input pattern, we compare the measured output and the predicted output computed with the measured TM. Then, we show that the distribution of the singular values of the measured TM varies according to random matrix theory. Finally, we demonstrate that single- or multi-point light focusing can be achieved. Interestingly, our Bayesian formulation allows the use of the same model for the estimation of the optimal DMD binary input pattern. In addition to being an interesting signal processing problem, this approach is particularly relevant for real-life applications of the TM approach, since it allows a simple, fast and robust implementation.




Fig. 1. Different experimental approaches for measuring the complex-valued transmission matrix of a scattering medium with a binary DMD amplitude modulator: (a) using a reference arm for retrieving the phase of the output field by off-axis or phase-shifting holography; (b) using the DMD as a spatial phase modulator by displaying amplitude holograms, and using the unmodulated parts of the field as a phase-stable reference; (c) presented approach where only intensity values are measured.

2. Experimental setup

Our experimental setup, described in Fig. 2, uses a DMD-array from Texas Instrument (1920×1080 tilting micromirrors), driven by the DLP V-9500 VIS module (Vialux). The DMD is made of mirrors that can switch between two angular positions separated by 24° , thus reflecting each pixel either toward a beam dump (pixel “OFF”) or towards the focusing system (pixel “ON”).

Under Matlab, an amplitude mask is computed and loaded on the DMD. The pattern corresponding to the ON pixels is focused on the surface of a thick scattering medium by means of a $f = 100$ mm lens L1 (thus the DMD pixels correspond roughly to incidence angles on the sample). The sample is a ~ 100 microns thick layer of white paint, which is thick enough in order to considerably mix the light on the other side, producing a complex speckle interfering pattern. This speckle pattern is collected through a microscope objective (L2) and detected on a camera (AVT Pike F-100B). In order to measure the TM, we need to send a large series of input patterns (typically a few times the number of input pixels we wish to control), in a time over which the medium can be considered stationary. For this purpose, we use the “high speed” driver provided with the DMD in order to load all the to-be-projected random amplitude masks to the memory of the DMD driver module, and we trigger the display of each mask via a DAQ card (National Instruments, PCI-6221) and a waveform generator. In the same way, in order to be as fast as possible, we also only consider a subregion on the camera of size 400×400 pixels. The overall acquisition rate is 31 images per second. To monitor the stability of the medium, we periodically measure the correlation of the speckle image corresponding to the same input mask. We therefore quantify the stability of the medium, which is better than 98%

over the total measurement time (typically around 5 minutes).




Fig. 2. Experimental scheme: A 532 nm CW laser is expanded through a telescope in order to obtain an homogeneous beam. Through a rectangular mask, it illuminates the DMD which acts as binary amplitude spatial light modulator. The DMD reflects the light in two different directions corresponding to either “ON” (unit transmission) or “OFF” (the light is deviated towards a beam dump). The transmitted pattern is focused by a first lens L1 on the scattering medium – here a white paint layer –, acting as a thick multiply scattering medium. The transmitted speckle pattern is collected by a microscope objective and is observed through a polarizer P on a CCD camera.

3. Phase retrieval for optics through complex media

The experimental setup described in the previous section presents different challenges in order to exploit the TM framework: the absence of reference requires a solution exploiting only intensity measurements, and this must be robust to experimental noise. This paper shows that signal processing techniques for phase retrieval can be used, first for the estimation of the TM, and then for getting the optimal DMD configuration in a focalization task.

3.1. Calibration as a phase retrieval problem

If the focalization task - or in a more general view, any front-shaping task using spatial light modulators - can be easily expressed as a phase retrieval problem - knowing the TM, which SLM setup best explains the intensity observations or the desired output -, this is not the case for the estimation of the TM.

We formalize the latter as a calibration problem: given P incoming waves, assumed perfectly known, which model explains at best the observed outputs?

Formally, let $\mathbf{x}_\mu \in \{0, 1\}^N$ stand for the binary DMD inputs related to the μ -th acquisition, where N is the number of pixels (mirrors) used on the DMD. We assume that the partial observations of the sole moduli of the transmitted waves (the square root of the camera measured intensities), denoted by $\mathbf{y}_\mu \in \mathbb{R}_+^M$, obey

$$\mathbf{y}_\mu = |\mathbf{D}\mathbf{x}_\mu|, \quad \forall \mu \in \{1, \dots, P\}, \quad (1)$$

where \mathbf{D} is the *unknown* complex-valued transmission matrix characterizing the scattering material, and M is the number of observed pixels on the camera.

Then, adopting a matrix formulation and conjugating-transposing the system, we get

$$\mathbf{Y}^H = |\mathbf{X}^H \mathbf{D}^H|, \quad (2)$$

where $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_P]$, $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_P]$ and \cdot^H denotes the conjugate-transpose of a matrix/vector. This reveals a “classic” phase retrieval problem: given the matrix of inputs \mathbf{X}^H , each column of \mathbf{Y}^H is used to estimate each complex-valued column of \mathbf{D}^H .

3.2. Phase retrieval techniques

The problem of reconstructing a complex vector given only the magnitude of measurements is a non-convex optimization problem notoriously difficult to solve. Many algorithms have been devised in the literature to deal with this problem, especially in optics where most approaches aim at reconstructing a phase object from an intensity image in a Fourier plane [24, 25]. Here we are interested in a somewhat simpler and more general problem where the transform is completely general and the output points can be treated independently. Amongst recent contributions to this signal processing problem, we can mention [26, 27] which approximate the phase recovery problem by relaxed problems enabling the use of standard optimization procedures, and the Bayesian procedures proposed in [23, 28], which circumvent the non-linearity of the modulus through the introduction of hidden variables and resort to variational approximations.

In this paper, we chose to use the so-called *prVBEM* Bayesian approach introduced in [23]. Our choice is motivated by two of its main advantages:

- its generic framework, enabling its use in both applications we are here interested in, namely the calibration of the TM and the focalization task,
- its reasonable computational complexity, of crucial importance for a complete procedure including the calibration of the TM and its immediate use for focusing.

We will not detail here the derivation of the algorithm. Its particularizations to the two problems of interest, calibration of the TM and focalization, are described resp. in appendix 7.1 and 7.2. For a methodological justification of the approach, we refer the reader to the original paper [23].

4. Estimating the TM with intensity-only measurements and binary inputs

In this section, we focus on the calibration stage of the proposed approach. We assess the accuracy of the TM estimated through the phase retrieval technique [23] with regard to two different analyses: its prediction performance, namely its ability to predict the observations, and its nature, expected to satisfy random features.

4.1. Prediction performance

To assess the prediction performance of the estimated TM, we adopt a cross-validation-like experimental framework. The setup is as follows. We measure the $M = 40000$ camera pixels stemming from $N = 900$ DMD mirrors, 50% of them being turned on, the others off at each displayed pattern. The operation is repeated randomly $P = 6000$ times. Given this dataset, a row of the TM is then learned from $p = \alpha N$ calibration measurements, with α varying in $\{1, \dots, 6\}$, and used in a second step to predict the $P - p$ remaining measurements. This estimation is performed on 50 different rows of the TM.

The *prVBEM* algorithm presents of complexity of order $\mathcal{O}(p^2)$. Run for 200 iterations, this leads to a computational time of order 0.6 s per estimated row of the TM, in Matlab, on a MacBook Air with a 1.7Ghz i7 processor. This performance is good in comparison with other state-of-the-art algorithms (see *e.g.*, [23]) and, keeping in mind that rows are independent, sufficient

for making possible a complete procedure “calibration of the TM-focalization” in a reasonable time.




Fig. 3. Prediction performance according to (a), the mean-square error (MSE), in log scale, and to (b), the normalized cross-correlation between observation predictions using the estimated TM, and actual measurements of the output moduli (square root of the camera intensity values), as a function of the number of calibration measurements (x-axis is α , such that $p = \alpha N$ calibration measurements are used).

The prediction performance of the estimated TM is evaluated according to the mean-square error (MSE) and the normalized cross-correlation between the moduli of the $P - p$ predicted measurements and the actual observed ones, such as, for the estimate $\hat{\mathbf{d}}$ of the current (conjugate) row of the TM (see Eq. (2)),

$$MSE = \frac{\|\mathbf{y}_{P-p} - \mathbf{X}_{P-p}^H \hat{\mathbf{d}}\|_2^2}{\|\mathbf{y}_{P-p}\|_2^2}, \quad (3)$$

$$Corr = \frac{\mathbf{y}_{P-p}^H \mathbf{X}_{P-p}^H \hat{\mathbf{d}}}{\|\mathbf{y}_{P-p}^H \mathbf{y}_{P-p}\|_2 \|\hat{\mathbf{d}}^H \mathbf{X}_{P-p}^H \mathbf{X}_{P-p}^H \hat{\mathbf{d}}\|_2}, \quad (4)$$

where \mathbf{y}_{P-p} stands for the current (real-valued) row of \mathbf{Y} (see Eq. (2)) restricted to its last $P - p$ elements, the first p ones being used for the estimation of $\hat{\mathbf{d}}$.

These figures of merit give different insights into the estimations achieved by the algorithm: the MSE measures the distance (in an euclidian sense) to the actual observations, while the correlation evaluates their angular difference. Figure 3 show these quantities (resp. in (a) and (b)) averaged over the 50 rows of the TM considered for estimation (red lines), as well as the minimum and maximum (given by whiskers), for an increasing $\alpha = p/N$.

Both curves match admirably: increasing α leads to a decrease of the MSE and an increase of the correlation. Interestingly, we see that for $\alpha \geq 3$, that is, for at least 3 times more real measurements than complex unknowns, the estimation of the TM is accurate enough to predict the observations with an average correlation around 0.95 and an average MSE lower than 0.18 (Fig. 3(a) shows this value in log-scale).

4.2. Comparison of singular values to Random matrix theory

Interestingly, we can check that the measured TM presents some characteristics as predicted by random matrix theory. One practical way is to verify that the distribution of its normalized




Fig. 4. Density of the normalized singular values for different $\gamma = M/N$. Stamped line: experimental results, continuous line: Marčenko-Pastur law.

singular values obeys the Marčenko-Pastur law [29]. It should be noted that such apparently random signals are the hardest case for phase retrieval, where no specific structure can be taken into account.

In order to reduce the influence of specifics of our experimental setting, we perform the following operations, as in [30]:

- i)* We normalize over the rows and columns, to attenuate the illumination artifacts: residual illumination “by default” on each pixel of the camera for the rows, and inhomogeneous contribution of each DMD mirror on the entire set of camera pixels for the columns.
- ii)* Because of the size of the speckle grains, two neighboring DMD mirrors may affect the material in the same way, as well, two pixels of the camera will be potentially correlated. To avoid this effect, we subsample the rows and columns of the matrix.

To draw the empirical spectral density, we then consider the following setup. We subsample the columns of the matrix up to $N = 200$ and leave the number of rows varying, more precisely $M = \gamma N$, with $\gamma \in \{1, \dots, 6\}$. These sub-matrices thus constitute partitions of the estimated matrix, randomly picked 100 times to average the resulting densities. Figure 4 compares the experimental curves to the theoretical ones drawn according to the Marčenko-Pastur law. We see that the experiments qualitatively follow the predictions. We remark however that the larger γ is, the more chances we have to consider the contributions of neighboring correlated pixels. This partly explains the increasing gap between both curves.

5. Focusing with the DMD

Knowing the TM gives a powerful and flexible tool to control light within the scattering medium [30]. In particular, it can be used to compute which DMD input has to be set, in order to display a given arbitrary pattern at the receiver end. In this section, we demonstrate the special case of

focusing light with maximum intensity on a desired pattern (a chosen sparse subset of the output pixels), with the TM measured experimentally as in the section 4. It should be emphasized that we keep the same experimental setup, with the binary DMD as input device. Here, simple inversion methods such as [30] cannot be used, as these require phase-modulated input and amplitude and phase detection.

We propose here to resort to the Bayesian variational approach [23], in a similar way as for the calibration but adapted here to the binary nature of the DMD inputs. The particularization of the algorithm to this case is given in appendix 7.2.




Fig. 5. Illustration of light focusing on 3 points. The circles mark the positions of the targets.

We assess the performance of the proposed focusing approach through different experiments. The general setting is as follows. The DMD inputs, here taken of dimension $N = 1600$, are estimated from the desired outputs (see appendix 7.2), focusing on 1 to 4 target points. The procedure exploits the TM, reduced to its rows of interest and previously measured as discussed in section 4.

Figure 5 shows an example of the observed output field, corresponding to the estimated DMD configuration, optimized to focus on 3 points. To quantitatively evaluate the focusing performance, we measure the intensity enhancement factor, as:

$$\eta \triangleq \frac{I_{\text{foc}}}{I_{\text{back}}}, \quad (5)$$

where I_{foc} is the intensity inside the target area after spatial binary amplitude modulation is performed, I_{back} is the average background intensity. This value is measured for 100 trials, as a function of the number of calibration measurements used to learn the TM.

Two different setups are then considered: the single-point focusing case and the multi-target case.

5.1. Focusing on a single point

Figure 6 compares the enhancement factors achieved by two different focusing methods, namely a simple binary phase-conjugation - performing $\hat{\mathbf{x}} = [\Re(\mathbf{D}^H \mathbf{y}) > 0]$ - and the proposed method, in the case where only one target point is focused. Results are presented under a “box” format, where:




Fig. 6. Single target experiment. Enhancement factor as a function of the number of measurements used to learn the TM (x-axis is α , such that $p = \alpha N$ calibration measurements are used). For the same estimation of the TM, 2 focusing techniques are compared: binary phase conjugation (blue boxes), and the Bayesian technique [23] (red boxes).

- the middle segment stands for the average enhancement $\bar{\eta}$ over the 100 trials,
- the upper and lower bounds of the rectangle define the interval $[\bar{\eta} - \sigma_{\eta} \quad \bar{\eta} + \sigma_{\eta}]$ (where σ_{η} is the experimentally computed standard deviation), in which lies, under the Gaussian assumption, 68 % of the trials,
- the whiskers represent the minimum and maximum values observed over the entire set of trials.

For each experiment point $\alpha \in \{2, \dots, 6\}$, such that $p = \alpha N$ calibration measurements are used to compute the TM, we display the boxes related to the binary phase-conjugation method (blue boxes), and the Bayesian technique (red boxes) [23] particularized in appendix 7.2.

As a first observation, we can see that the general dependency with regard to α noticeably resonates with the curve of the *prVBEM* algorithm in Fig. 3(b): there is a clear gap between the performance achieved for $\alpha = 2$ and for $\alpha = 3$, while, for $\alpha \geq 3$, the intensity enhancement keeps increasing but less significantly.

Interestingly, the *prVBEM* algorithm seems to outperform binary phase-conjugation, with regard to the mean and maximum values measured, but not in a statistically significant manner. Focusing on the most favorable case considered here, namely with $\alpha = 6$, the best intensity enhancement factor lies around 140, to be compared with the ideal expected enhancement given by $1 + \frac{1}{\pi} \left(\frac{N}{2} - 1 \right) \simeq 255$, see [17].

5.2. Focusing on multiple points

For this second setup, we are interested in the performance of the *prVBEM* approach in a context of multiple target points. Additionally to the intensity enhancement, we consider here the missed detection rate, defined as the number of trials (expressed in percentage) failing to focus on *at least* one of the multiple target points, *i.e.*, the number of trials for which at least one of the T largest intensity peaks in the output image does not match any of the T targets.




Fig. 7. Multiple target experiment. (a) Average enhancement factor as a function of the number of measurements used to learn the TM (x-axis is α , such that $p = \alpha N$ calibration measurements are used), and the number of target points (y-axis). (b) Missed detection rate (same axis as in (a)).

Figure 7 represents these two figures of merit under diagram formats. They present an interesting general symmetry: increasing the number of targets or decreasing the number of calibration points leads to an increase of the missed detections and a decrease of the enhancement factor. The missed detection rate seems however more sensitive to the number of calibration points used to learn the TM: for $\alpha = 2$ and 2 target points, the algorithm fails with a rate approaching 40%, while for $\alpha = 3$ and the same number of targets, we keep a reasonable performance (around 10%). In a more general view, these figures greatly highlight the deep relation between the quality of the calibration and the focusing performance.

6. Conclusion

This paper shows that the full *complex-valued* transmission matrix of a strongly scattering material can be estimated, up to a global phase factor on each of its rows, with a simple experimental setup involving only *real-valued* inputs and outputs. In our experiment, the inputs are amplitude modulations on a *binary* DMD, and the output is the field intensity measured on a CCD camera, that gathers a significant amount of measurement noise. Note that no reference arm is used, that would allow interferometric measurements, but that would make the experimental setup more complex and considerably more unstable.

We here resort to Bayesian phase retrieval techniques, and we have shown that, amongst such techniques, a recently proposed variational approach (the so-called *prVBEM* algorithm [23]) allows a precise estimation of the transmission matrix, tractable in computational complexity and scalable for large-size signals, provided that we have a sufficiently large number of input-output calibration signals. Experimental results validate this concept, both in terms of output prediction, distribution of singular values, and in an application of light focusing onto a number of target points in the output plane. It should be emphasized that this estimation of the transmission matrix opens many applications beyond light focusing, may it be for imaging through the scattering material [30, 31], or for obtaining information about the scattering material itself.

7. Appendix: particularization of the *prVBEM* algorithm

In this appendix, we outline the main lines of the *prVBEM* algorithm particularized to its use for calibration of the TM and focalization. We refer the reader to the original paper [23] for a methodological presentation of the procedure.

7.1. *prVBEM for the calibration of the TM*

To circumvent the non-linearity induced by the absence of phase observations, the procedure introduces new variables modeling, on the one hand, the missing phases of the observations, and on the other hand, some acquisition noise. Thus, recalling that we resort to a conjugate-transposition of the matrix system (see Eq. (2)), each absolute-valued measurement y_μ , $\mu \in \{1 \dots P\}$, of any row \mathbf{y} of \mathbf{Y} , is expressed as

$$y_\mu = e^{j\theta_\mu} \left(\sum_{i=1}^N x_{\mu i} d_i^* + \omega_\mu \right), \quad (6)$$

where $\theta_\mu \in [0, 2\pi)$ stands for its missing conjugate phase, $x_{\mu i}$ is the i th element of the μ th row in \mathbf{X} , d_i^* corresponds to the i th conjugate element in the current estimated row \mathbf{d} of \mathbf{D} and ω_μ is an additive noise, assumed centered isotropic Gaussian (denoted \mathcal{CN} in the following) with variance σ^2 . We moreover suppose that the probability distributions for the entries of the matrix and for the missing phases are:

$$p(\mathbf{d}) = \prod_{i=1}^N p(d_i) \quad \text{with} \quad p(d_i) = \mathcal{CN}(0, \sigma_d^2), \quad (7)$$

$$\text{and} \quad p(\theta) = \prod_{\mu=1}^P p(\theta_\mu) \quad \text{with} \quad p(\theta_\mu) = \frac{1}{2\pi}. \quad (8)$$

Under these assumptions, the absence of phases in the observations is naturally taken into account in the model since marginalizing on θ_μ leads to a distribution on y_μ which only depends on the moduli of y_μ and $\sum_{i=1}^N x_{\mu i} d_i^*$.

Within model (6)-(8), the recovery of the complex row \mathbf{d} of the TM can be expressed as the solution of the following marginalized Maximum A Posteriori (MAP) estimation problem

$$\hat{\mathbf{d}} = \underset{\mathbf{d}}{\operatorname{argmax}} p(\mathbf{d}|\mathbf{y}), \quad (9)$$

$$\text{with} \quad p(\mathbf{d}|\mathbf{y}) = \int_{\theta} p(\mathbf{d}, \theta|\mathbf{y}). \quad (10)$$

Because of the marginalization on the hidden variables θ , the direct computation of $p(\mathbf{d}|\mathbf{y})$ is however intractable in general. The *prVBEM* algorithm is based on the computation of a particular (the so-called ‘‘mean-field’’) approximation $\hat{q}(\mathbf{d}, \theta) = \prod_i q(d_i) \prod_\mu q(\theta_\mu)$ of the posterior joint distribution $p(\mathbf{d}, \theta|\mathbf{y})$. The procedure is then iterative, updating the factors $q(d_i)$ as

$$q(d_i) = \mathcal{CN}(m_i, \Sigma_i), \quad (11)$$

with

$$m_i = \frac{\sigma_d^2}{\sigma_n^2 + \sigma_d^2 \mathbf{x}_i^H \mathbf{x}_i} \langle \mathbf{r}_i \rangle^H \mathbf{x}_i, \quad \Sigma_i = \frac{\sigma_n^2 \sigma_d^2}{\sigma_n^2 + \sigma_d^2 \mathbf{x}_i^H \mathbf{x}_i}, \quad \langle \mathbf{r}_i \rangle = \bar{\mathbf{y}} - \sum_{k \neq i} m_k^* \mathbf{x}_k, \quad (12)$$

and

$$\bar{\mathbf{y}} = \left[y_\mu e^{j \arg(y_\mu^* \langle z_\mu \rangle)} \frac{I_1\left(\frac{2}{\sigma_n^2} |y_\mu^* \langle z_\mu \rangle|\right)}{I_0\left(\frac{2}{\sigma_n^2} |y_\mu^* \langle z_\mu \rangle|\right)} \right]_{\mu=\{1 \dots P\}} \quad (13)$$

$$\langle z_\mu \rangle = \sum_{i=1}^N m_i^* x_{\mu i}. \quad (14)$$

In the equations above, I_0 (resp. I_1) stands for the modified Bessel function of the first kind for order 0 (resp. 1).

Coming back to problem (9), an approximation of $p(\mathbf{d}|\mathbf{y})$ simply follows from

$$p(\mathbf{d}|\mathbf{y}) = \int_{\theta} p(\mathbf{d}, \theta | \mathbf{y}), \quad (15)$$

$$\simeq \int_{\theta} q(\mathbf{d}, \theta), \quad (16)$$

$$\simeq \int_{\theta} \prod_i q(d_i) \prod_{\mu} q(\theta_{\mu}), \quad (17)$$

$$\simeq \prod_i q(d_i), \quad (18)$$

resulting in an element-wise estimation of row $\hat{\mathbf{d}}$, such as $\hat{d}_i = m_i$.

7.2. prVBEM for focusing

The problem of focusing is expressed as an inverse problem, where, knowing the TM \mathbf{D} and the observation \mathbf{y} , we look for the DMD input \mathbf{x} such as described in (1). Adopting a similar modeling as in previous section, we then assume, for all elements y_μ with $\mu \in \{1, \dots, M\}$,

$$y_\mu = e^{j\theta_\mu} \left(\sum_{i=1}^N d_{\mu i} x_i + \omega_\mu \right), \quad (19)$$

where $\theta_\mu \in [0, 2\pi)$ stands for the missing conjugate phase, $d_{\mu i}$ is the μ th element of the i th column in \mathbf{D} , $x_i \in \{0, 1\}$ corresponds to the state of the i th DMD pixel and ω_μ is an additive noise, assumed centered isotropic Gaussian of variance σ^2 . As previously, we suppose that the elements θ_μ are independently and uniformly distributed in the interval $[0, 2\pi)$, however, in order to accommodate for binary inputs, we consider here a Bernoulli model for \mathbf{x} :

$$p(\mathbf{x}) = \prod_{i=1}^N p(x_i) \quad \text{with} \quad p(x_i) = \text{Ber}(p_i) = \begin{cases} p_i & \text{if } x_i = 1, \\ 1 - p_i & \text{if } x_i = 0, \end{cases} \quad (20)$$

where we set p_i to 0.5, noticing that asymptotically half of the DMD pixels are expected to be ‘‘ON’’ [17].

Then, within model (19)-(20), the recovery of the DMD inputs is expressed as the marginalized MAP estimation problem

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\text{argmax}} p(\mathbf{x}|\mathbf{y}), \quad (21)$$

$$\text{with} \quad p(\mathbf{x}|\mathbf{y}) = \int_{\theta} p(\mathbf{x}, \theta | \mathbf{y}). \quad (22)$$

The procedure is then similar to the one described previously for the calibration task. The *prVBEM* algorithm iteratively updates the factors of the mean-field approximation $q(\mathbf{x}, \boldsymbol{\theta}) = \prod_i q(x_i) \prod_\mu q(\boldsymbol{\theta}_\mu)$ of $p(\mathbf{x}, \boldsymbol{\theta} | \mathbf{y})$, which leads, in the particular case of focusing, to the expression

$$q(x_i) = p(x_i) \exp\left(x_i \frac{2 \Re(\mathbf{d}_i^H \langle \mathbf{r}_i \rangle) - \mathbf{d}_i^H \mathbf{d}_i}{\sigma^2}\right), \quad (23)$$

where

$$\langle \mathbf{r}_i \rangle = \bar{\mathbf{y}} - \sum_{k \neq i} q(x_k = 1) \mathbf{d}_k, \quad (24)$$

$$\bar{\mathbf{y}} = \left[y_\mu e^{(j \arg(y_\mu^* \langle z_\mu \rangle))} \frac{I_1\left(\frac{2}{\sigma^2} |y_\mu^* \langle z_\mu \rangle|\right)}{I_0\left(\frac{2}{\sigma^2} |y_\mu^* \langle z_\mu \rangle|\right)} \right]_{\mu=\{1 \dots M\}} \quad (25)$$

$$\langle z_\mu \rangle = \sum_i q(x_i = 1) d_{\mu i}, \quad (26)$$

and I_0 (resp. I_1) stands for the modified Bessel function of the first kind for order 0 (resp. 1).

Then, an approximation of $p(\mathbf{x} | \mathbf{y})$ is given by

$$p(\mathbf{x} | \mathbf{y}) = \int_{\boldsymbol{\theta}} p(\mathbf{x}, \boldsymbol{\theta} | \mathbf{y}), \quad (27)$$

$$\simeq \int_{\boldsymbol{\theta}} \prod_i q(x_i) \prod_\mu q(\boldsymbol{\theta}_\mu), \quad (28)$$

$$= \prod_i q(x_i). \quad (29)$$

Using this approximation, the problem is easy to solve by a simple thresholding operation, *i.e.*, $\hat{x}_i = 1$ if $q(x_i = 1) > 0.5$ and $\hat{x}_i = 0$ otherwise.

Acknowledgements

AD is currently working at ENSTA Bretagne, LabSTICC (UMR 6285), 2 rue François Verny, F-29200 Brest, France. OK acknowledges the support of the Marie Curie Intra-European Fellowship for career development (IEF). LD acknowledges a joint research position with the Institut Universitaire de France. This work has been supported in part by the CSI:PSL grant and LABEX WIFI (Laboratory of Excellence within the French Program “Investments for the Future”) under references ANR-10-LABX-24 and ANR-10-IDEX-0001-02 PSL*, and by the ERC under the European Union’s 7th Framework Programme Grant Agreements 307087-SPARCS and 278025-COMEDIA.

DOA ESTIMATION IN STRUCTURED PHASE-NOISY ENVIRONMENTS

Angélique Drémeau^{⊗*}, Cédric Herzet^{*}

[⊗] ENSTA Bretagne and Lab-STICC UMR 6285, Brest, F-29200, France

^{*} INRIA Centre Rennes-Bretagne Atlantique, Campus universitaire de Beaulieu, F-35000 Rennes, France

ABSTRACT

In this paper we focus on the problem of estimating the directions of arrival (DOA) of a set of incident plane waves. Unlike many previous works, which assume that the received observations are only affected by additive noise, we consider the setup where some phase noise also corrupts the data (as for example observed in atmospheric sound propagation or underwater acoustics). We propose a new methodology to solve this problem in a Bayesian framework by resorting to a variational mean-field approximation. Our simulation results illustrate the benefits of carefully accounting for the phase noise in the DOA estimation process.

Index Terms— Direction-of-arrival estimation, phase noise, variational Bayesian approximation, mean-field approximation

1. INTRODUCTION

Estimating the directions of arrival (DOA) of propagating plane waves is at the heart of many applicative domains including sonar, radar and mobile telecommunications. Among the rich literature dealing with this problem, the most popular method is probably conventional beamforming [1] which can roughly be interpreted as a least-square estimator. *Per se*, this approach suffers from a lower limit on the resolution achievable in the DOA estimation process (conditioned by the length of the sensor array). To overcome this issue, so-called “high-resolution” techniques, taking advantage of more prior information on (the number and the nature of) the sources, have been proposed in the literature (see *e.g.*, [2, 3, 4, 5]). In [2], the authors introduced the well-known MUSIC algorithm which benefits from the knowledge of the number of the sources and rely on the assumption that the noise and the signal of interest live in perfectly separable subspaces. More recently, a “compressive” beamforming approach was proposed in [3, 4, 5], where a sparse prior was exploited to address the DOA estimation problem.

The contributions mentioned above assume that the incident plane waves are only corrupted by some additive noise. Unfortunately, when the waves travel through highly fluctuating media, as in the case of *e.g.*, atmospheric sound

propagation [6] or underwater acoustics [7], this model does no longer describe accurately the physics underlying the propagation process. In such cases, a multiplicative phase noise typically corrupts the collected signal, making the corresponding DOA estimation problem quite challenging. We address this problem in the present paper.

Our approach is inspired from the recent standard “high-resolution” DOA methods [3, 4, 5] and some phase retrieval algorithms presented in [8, 9]. More specifically, we model the received signal as a sparse combination of elementary signals (taken from a redundant dictionary) and assume that the latter is corrupted by both additive and phase noise. Our methodology is grounded on a probabilistic Bayesian framework and relies on a variational mean-field approximation. In particular, we show how to nicely incorporate fine noise-phase models in this framework, extending in this respect, the approaches proposed in [8, 9].

2. PROBLEM STATEMENT

Our derivations are based on the following formulation of the DOA estimation problem: we consider an antenna composed of N sensors and assume that the collected observation vector $\mathbf{y} \in \mathbb{C}^N$ can be expressed as

$$\mathbf{y} = \mathbf{P}\mathbf{D}\mathbf{z} + \boldsymbol{\omega}, \quad (1)$$

where $\boldsymbol{\omega} \in \mathbb{C}^N$ and $\mathbf{P} = \text{diag}(\{e^{j\theta_n}\}_{n=1}^N) \in \mathbb{C}^{N \times N}$ play respectively the role of an additive and a multiplicative phase noise. Matrix $\mathbf{D} = [\mathbf{d}_1 \dots \mathbf{d}_M] \in \mathbb{C}^{N \times M}$ is made up of the steering vectors $\mathbf{d}_i \triangleq [e^{j\frac{2\pi}{\lambda}\Delta \sin(\phi_i)} \dots e^{j\frac{2\pi}{\lambda}\Delta N \sin(\phi_i)}]^T$, where ϕ_i 's are some potential angles of arrival, Δ is the distance between two adjacent sensors, and λ is the wavelength of the propagation waves.

With this formulation, assuming that \mathbf{y} results from the combination of a few waves arriving from different angles ϕ_i 's, the DOA estimation problem is basically equivalent to identifying the positions of the nonzero coefficients in \mathbf{z} (since each column of \mathbf{D} corresponds to a particular angle of arrival). In the standard DOA estimation problem, \mathbf{P} is assumed to be known with $\mathbf{P} = \mathbf{I}_N$, where \mathbf{I}_N is the $N \times N$ identity matrix. In this case, the model connecting the unknown vector \mathbf{z} to the measurements \mathbf{y} is linear; finding the

^{*}This work has been supported by the DGA/MRIS.

position of the nonzero elements in \mathbf{z} can then be carried out with standard sparse-representation algorithms, see *e.g.*, [10].

In this paper, we consider the more complex case where $\mathbf{P} = \text{diag}(\{e^{j\theta_n}\}_{n=1}^N)$ and the θ_n 's are unknown. More specifically, we assume that the phases θ_n 's obey the following Markov model:

$$p(\boldsymbol{\theta}) = \prod_{n=2}^N p(\theta_n|\theta_{n-1}) p(\theta_1), \quad (2)$$

with $p(\theta_n|\theta_{n-1}) = \mathcal{N}(a\theta_{n-1}, \sigma_\theta^2)$, $\forall n \in \{2, \dots, M\}$, $a \in \mathbb{R}_+$, and $p(\theta_1) = \mathcal{N}(0, \sigma_1^2)$. From a practical point of view, this model allows us to describe spatial fluctuations of the propagation medium all along the antenna; the strength of the fluctuations is related to the value of parameter σ_θ^2 .

As noted in the introduction, assuming that \mathbf{P} is unknown renders the DOA estimation much more difficult since it introduces uncertainties in the observation model. Before proceeding to the presentation of the proposed methodology to address this issue, we draw some connections with other applicative fields:

- Our work relates first to the phase retrieval problem (*e.g.*, [11]) where the phase information of the observations is completely missing: only intensities or amplitudes are acquired. Formally, both problems share - explicitly or not - the same observation model (1) but differ in the prior distribution they enforce on the phases $\boldsymbol{\theta}$, the absence of phase information being modeled by a non-informative prior, such as a uniform law (see [8, 9]).
- Then we note that, in top of being relevant for the DOA estimation problem with fluctuating media, this model is also of interest in the domain of digital communications where it can be used to characterize the transmission of complex modulation symbols over a channel affected by carrier phase noise, see *e.g.*, [12].

3. BAYESIAN FORMULATION OF THE PROBLEM

We address the problem of estimating \mathbf{z} from \mathbf{y} when the realizations of $\boldsymbol{\omega}$ and $\boldsymbol{\theta}$ are unknown. To that end, we first place this problem into a Bayesian framework by defining suitable additional prior distributions on the unknown quantities.

To account for the sparsity of \mathbf{z} , we suppose that, $\forall i \in \{1, \dots, M\}$, $z_i = x_i \cdot s_i$ with

$$p(x_i) = \mathcal{CN}(0, \sigma_x^2), \quad \text{and} \quad p(s_i) = \text{Ber}(p_i). \quad (3)$$

This so-called Bernoulli-Gaussian model has been now largely used in the literature to model sparse priors (see *e.g.*, [13]). Next, we assume that $p(\boldsymbol{\omega})$ is a zero-mean Gaussian distribution with variance σ^2 . This hypothesis is typically justified by the central limit theorem under the assumption

that the additive noise corrupting the data results from the aggregation of a large number of random parasitic contributions. Finally, the probabilistic model describing the behavior of $\boldsymbol{\theta}$ is given by (2). As mentioned in the previous section, this simple model accounts for local variations of the propagation medium along the sensor network.

Based on this probabilistic model, we propose to look for the solution of the following Minimum Mean Square Error (MMSE) problem

$$\hat{\mathbf{z}} = \arg \min_{\tilde{\mathbf{z}}} \mathbb{E}_{\mathbf{z}|\mathbf{y}} [\|\mathbf{z} - \tilde{\mathbf{z}}\|_2^2], \quad (4)$$

relying on the marginal posterior distribution $p(\mathbf{z}|\mathbf{y}) = \int_{\boldsymbol{\theta}} p(\mathbf{z}, \boldsymbol{\theta}|\mathbf{y}) d\boldsymbol{\theta}$. The computation of this marginalization being an intractable problem, we propose hereafter a practical procedure based on a mean-field approximation of the joint distribution $p(\mathbf{z}, \boldsymbol{\theta}|\mathbf{y})$ to approach the solution of (4).

4. THE PROPOSED PROCEDURE

Mean-field approximations aim at approximating a posterior joint distribution by another one constrained to have a ‘‘simple’’ factorization while minimizing some distance with the targeted distribution. In the following, we will look for an approximation of $p(\mathbf{z}, \boldsymbol{\theta}|\mathbf{y})$, say $\hat{q}(\mathbf{z}, \boldsymbol{\theta})$, obeying the following factorization $\hat{q}(\mathbf{z}, \boldsymbol{\theta}) = \hat{q}(\boldsymbol{\theta}) \prod_{i=1}^M \hat{q}(z_i)$. With this approximation of $p(\mathbf{z}, \boldsymbol{\theta}|\mathbf{y})$, the evaluation of the marginal with respect to \mathbf{z} is simplified since $\int_{\boldsymbol{\theta}} \hat{q}(\mathbf{z}, \boldsymbol{\theta}) = \prod_i \hat{q}(z_i)$ and the solution of the MMSE problem (4) can be approximated component-wise as

$$\hat{z}_i \simeq \int_{z_i} z_i \hat{q}(z_i) dz_i.$$

A well-known algorithm to find a mean-field approximation of a target distribution is the so-called ‘‘Variational Bayes Expectation-Maximization’’ (VBEM) algorithm (see *e.g.*, [14]). Particularized to our problem, this procedure searches for a local minimum of the following optimization problem:

$$\hat{q}(\mathbf{z}, \boldsymbol{\theta}) = \arg \min_q \int_{\mathbf{z}, \boldsymbol{\theta}} q(\mathbf{z}, \boldsymbol{\theta}) \log \left(\frac{q(\mathbf{z}, \boldsymbol{\theta})}{p(\mathbf{z}, \boldsymbol{\theta}|\mathbf{y})} \right) d\mathbf{z} d\boldsymbol{\theta}$$

$$\text{subject to } q(\mathbf{z}, \boldsymbol{\theta}) = q(\boldsymbol{\theta}) \prod_{i=1}^M q(z_i),$$

by sequentially minimizing the cost function with respect to $q(\boldsymbol{\theta})$ and $q(z_i)$, $\forall i \in \{1, \dots, M\}$.

We show hereafter that the sequence of distributions generated by the VBEM algorithm, say $\{q^{(k)}(\boldsymbol{\theta}), \{q^{(k)}(z_i)\}_i\}_k$, admit simple parametric expressions. We detail these expressions (at a given iteration of the procedure) in the following subsections. Due to space limitation, we do not provide the technical derivations but refer the reader to our companion report [15] for more details. For the sake of clarity, we also omit the iteration index (k) in the notations.

4.1. Update of $q(\theta)$

Under the condition¹ of small σ_θ^2 , we can express the estimate of $q(\theta)$ as a Gaussian distribution:

$$q(\theta) = \mathcal{N}(\mathbf{m}_\theta, \Sigma_\theta),$$

where
$$\Sigma_\theta^{-1} = \Lambda_\theta^{-1} + \text{diag} \left(\frac{2}{\sigma^2} |\boldsymbol{\eta}| \right),$$

$$\mathbf{m}_\theta = \Sigma_\theta \left(\text{diag} \left(\frac{2}{\sigma^2} |\boldsymbol{\eta}| \right) \arg(\boldsymbol{\eta}) \right),$$

and Λ_θ^{-1} is the precision matrix attached to the prior distribution on θ (2), *i.e.*,

$$\Lambda_\theta^{-1} = \begin{pmatrix} \frac{1}{\sigma_1^2} + \frac{a^2}{\sigma_\theta^2} & -\frac{a}{\sigma_\theta^2} & 0 & 0 \\ -\frac{a}{\sigma_\theta^2} & \frac{1+a^2}{\sigma_\theta^2} & \ddots & 0 \\ 0 & \ddots & \ddots & -\frac{a}{\sigma_\theta^2} \\ 0 & 0 & -\frac{a}{\sigma_\theta^2} & \frac{1}{\sigma_\theta^2} \end{pmatrix}. \quad (5)$$

Vector $\boldsymbol{\eta}$ is defined as $\boldsymbol{\eta} \triangleq [\eta_1, \dots, \eta_N]^T$, with

$$\eta_n = y_n \sum_i \langle z_i \rangle^* d_{ni}^*,$$

where d_{ni}^* is the conjugate of the n th element of \mathbf{d}_i , and $\langle z_i \rangle$ is defined using the current estimate of $q(z_i)$ as

$$\langle z_i \rangle \triangleq \int_{z_i} z_i q(z_i) dz_i.$$

Note that the distribution $q(\theta)$ being Gaussian, the marginals $q(\theta_n)$ come straightforwardly as

$$q(\theta_n) = \mathcal{N}(m_{\theta_n}, \Sigma_{\theta_n}),$$

where m_{θ_n} (resp. Σ_{θ_n}) is the n th element in \mathbf{m}_θ (resp. in the diagonal of Σ_θ). In practice, estimating these parameters can be efficiently implemented through a Kalman smoother [16] because of the particular structure of the precision matrix (5).

4.2. Update of $q(z_i)$

The iterates of the $q(z_i)$'s take the form of a mixture of two Gaussian distributions and only depend on the marginal distributions $q(\theta_n)$ of $q(\theta)$. More precisely, given that

$$q(z_i) = q(x_i, s_i) = q(x_i | s_i) q(s_i),$$

we obtain

$$q(x_i | s_i) = \mathcal{CN}(m_{x_i}(s_i), \Sigma_{x_i}(s_i)), \quad (6)$$

$$q(s_i) \propto \sqrt{\Sigma_{x_i}(s_i)} \exp \left(\frac{m_{x_i}(s_i)^* m_{x_i}(s_i)}{\Sigma_{x_i}(s_i)} \right) p(s_i), \quad (7)$$

¹We justify this condition in the technical report [15].

where \propto and \cdot^H means respectively proportionality and transpose complex conjugate,

$$\Sigma_{x_i}(s_i) = \frac{\sigma^2 \sigma_x^2}{\sigma^2 + s_i \sigma_x^2 \mathbf{d}_i^H \mathbf{d}_i}, \quad (8)$$

$$m_{x_i}(s_i) = s_i \frac{\sigma_x^2}{\sigma^2 + s_i \sigma_x^2 \mathbf{d}_i^H \mathbf{d}_i} \mathbf{d}_i^H \langle \mathbf{r}_i \rangle, \quad (9)$$

$$\langle \mathbf{r}_i \rangle = \bar{\mathbf{y}} - \sum_{k \neq i} q(s_k = 1) m_{x_k}(s_k = 1) \mathbf{d}_k, \quad (10)$$

$$\bar{\mathbf{y}} = \left[y_n e^{-j m_{\theta_n}} \frac{I_1(1/\Sigma_{\theta_n})}{I_0(1/\Sigma_{\theta_n})} \right]_{n=\{1 \dots M\}}, \quad (11)$$

and I_0 (resp. I_1) stands for the modified Bessel of the first kind of order 0 (resp. 1).

We note that the update equation (6)-(11) share some connections with the phase retrieval algorithm presented in [9]. The latter, relying also on a VBEM algorithm, differs from the proposed procedure in the definition of the prior distributions (2) and (3), respectively replaced by a uniform and a Gaussian distributions. In practice, both procedures share a similar structure. Leaving out the choice made here of a sparse-enforcing prior on \mathbf{z} , the main difference lies in the ‘‘reconstructed’’ phases m_{θ_n} in (11): while their definition relies here on the parameters of the Markov chain through the precision matrix (5), they only depend on the observations in [9] where a non-informative prior is considered.

4.3. Noise estimation

As emphasized in [13], the estimation of model parameters can easily be embedded within the VBEM procedure. Among them, the noise variance is of particular interest. Measure of the (mean) discrepancies between the observations and the assumed model, its iterative estimation usually helps the convergence of the algorithm to a proper local minimum, as observed in [13]. Particularized to model (1)-(3), this leads to

$$\begin{aligned} \hat{\sigma}^2 &= \arg \max_{\sigma^2} \int_{\mathbf{z}} q(\mathbf{z}, \boldsymbol{\theta}) \log p(\mathbf{z}, \boldsymbol{\theta}, \mathbf{y}; \sigma^2) d\mathbf{z} d\boldsymbol{\theta}, \\ &= \frac{1}{N} \left(\mathbf{y}^H \mathbf{y} - 2 \sum_i \Re \left[q(s_i = 1) m_{x_i}(s_i = 1) \bar{\mathbf{y}}^H \mathbf{d}_i \right] \right. \\ &\quad \left. + \sum_i \sum_{k \neq i} q(s_i = 1) q(s_k = 1) m_{x_i}(s_i = 1)^* m_{x_k}(s_k = 1) \mathbf{d}_i^H \mathbf{d}_k \right. \\ &\quad \left. + \sum_i q(s_i = 1) \left(\Sigma_{x_i}(s_i = 1) + |m_{x_i}(s_i = 1)|^2 \right) \mathbf{d}_i^H \mathbf{d}_i \right). \end{aligned}$$

In the following, we will refer to the proposed procedure as ‘‘paVBEM’’ for ‘‘phase-aware VBEM algorithm’’.

5. EXPERIMENTS

In this section, we assess numerically the effectiveness of the proposed approach. We consider the problem of identifying




Fig. 1. Evolution of the (averaged) normalized correlation as a function of the variance σ^2 when $K = 2$.

the directions of arrival of K plane waves from $N = 256$ observations (K will be specified later on). We assume that the angles of the K incident waves can be written as $\phi_k = -\frac{\pi}{2} + i_k \frac{\pi}{50}$ with $i_k \in [1, 50], \forall k \in \{1, \dots, K\}$. The set of angles $\{\phi_i = -\pi + i \frac{\pi}{50}\}_{i \in \{1, \dots, 50\}}$ together with the choice of the parameter $\Delta/\lambda = 4$ define the columns of the dictionary \mathbf{D} (see section 2). We set the following parameters for the phase Markov model (2): $\sigma_1^2 = 10^6$, $\sigma_\theta^2 = 1$ and $a = 0.8$. This corresponds to the situation where one has a large uncertainty on the initial value of the phase noise but connections exist between the phase noise on adjacent sensors.

As well as for the classic phase retrieval problem, the solution of (4) can only be found up to a global phase. Hence, computing the mean square distance between the ground truth \mathbf{z} and its reconstruction $\hat{\mathbf{z}}$ does not constitute a relevant figure of merit. Instead, we consider their normalized correlation, $\frac{|\hat{\mathbf{z}}^H \mathbf{z}|}{\|\hat{\mathbf{z}}\|_2 \|\mathbf{z}\|_2}$. This quantity is averaged over 50 realizations for each point of simulation. We compare the performance of the following algorithms: *i*) the conventional beamforming introduced in [1] ($\text{---}\circ\text{---}$); *ii*) the so-called *prVBEM* algorithm proposed in [9] as a solution to the phase retrieval problem ($\text{---}\triangle\text{---}$); *iii*) the *paVBEM* procedure described in section 4 ($\text{---}\diamond\text{---}$); *iv*) a relaxed version of *paVBEM* in which the sparsity of \mathbf{z} is not exploited but replaced by a Gaussian prior ($\text{---}\square\text{---}$).

The performance of these procedures are illustrated in Fig. 1 and 2 as a function of the noise variance σ^2 for $K = 2$ and $K = 5$, respectively. We see that the beamforming algorithm, which was originally proposed to solve the DOA estimation problem in the standard linear setup ($\mathbf{P} = \mathbf{I}_N$), fails to cope with the presence of fluctuations in the phase θ . The three other algorithms achieve different levels of performance, depending on the power of the additive noise and the number of incident waves. We note that all these procedures are derived from a similar optimization procedure




Fig. 2. Evolution of the (averaged) normalized correlation as a function of the variance σ^2 when $K = 5$.

but consider different degrees of knowledge on θ and \mathbf{z} . In [9], the authors assume that the phase is uniformly distributed and the sparse nature of \mathbf{z} is ignored; in the relaxed version of the proposed procedure (see [15]) the phase model (2) is exploited but the sparsity of \mathbf{z} is not taken into account; finally, as explained previously, the methodology presented in this paper integrates both the phase model (2) and the sparsity of \mathbf{z} in the estimation process. We see from Fig. 1 and 2 that the performance of these algorithms directly relates to the level of information they exploit: the proposed methodology outperforms its relaxed counterpart which, in turn, leads to better performance than the procedure proposed in [9].

We also notice that the procedures achieve better performance when $K = 5$ than $K = 2$. This counter-intuitive behavior is typical for phase-retrieval problems. In fact, it is easy to see that, when no phase information is available (*i.e.*, θ_n is uniform on $[0, 2\pi]$), only the modulus of y_n provides some information on \mathbf{z} . In such a case, the worst situation occurs when $K = 1$ since all the elements of \mathbf{D} has the same modulus (equal to 1) and the observations thus provide no information on \mathbf{z} . The setup $K = 2$ is close to this worst case, hence explaining the observed behavior.

6. CONCLUSION

We have presented a novel algorithm able to estimate DOA in environments corrupted by phase noise. Our approach relies on a mean-field approximation and exploits two types of priors: on the DOA through a sparse-enforcing distribution and on the phase noise through a Markov model. Our experiments have confronted the proposed approach to conventional beamforming and similar variational approaches handicapped by non-informative priors. In this regard, its good performance tends to prove a successful inclusion of the priors. Future work will include further assessment in underwater acoustics.

7. REFERENCES

- [1] D. H. Johnson and D. E. Dudgeon, *Array signal processing: concepts and techniques*, Englewood Cliffs, NJ, 1993.
- [2] R. O. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Trans. Antennas Propagation*, vol. AP-34, pp. 276–280, 1986.
- [3] W. Mantzel, J. Romberg, and K. Sabra, "Compressive matched-field processing," *Journal of Acoustical Society of America*, vol. 132, no. 1, pp. 90–102, 2012.
- [4] A. Xenaki, P. Gerstoft, and K. Mosegaard, "Compressive beamforming," *Journal of Acoustical Society of America*, vol. 136, no. 1, pp. 260–271, 2014.
- [5] S. Fortunati, R. Grasso, F. Gini, M. S. Greco, and K. Lepage, "Single snapshot DOA estimation by using compressed sensing," *EURASIP Journal on Advances in Signal Processing*, vol. 120, 2014.
- [6] S. Cheinet, L. Ehrhardt, D. Juvé, and P. Blanc-Benon, "Unified modeling of turbulence effects on sound propagation," *Journal of Acoustical Society of America*, vol. 132, no. 4, pp. 2198–2209, 2012.
- [7] R. Dashen, S.M. Flatté, W.H. Munk, K.M. Watson, and F. Zachariassen, *Sound Transmission Through a Fluctuating Ocean*, Cambridge Monographs on Mechanics. Cambridge University Press, 2010.
- [8] P. Schniter and S. Rangan, "Compressive phase retrieval via generalized approximate message passing," in *Communication, Control, and Computing (Allerton)*, October 2012.
- [9] A. Drémeau and F. Krzakala, "Phase recovery from a bayesian point of view: the variational approach," in *Proc. IEEE Int'l Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Brisbane, Australia, April 2015, pp. 3661–3665.
- [10] S. Foucart and H. Rauhut, *A mathematical introduction to compressive sensing.*, Applied and Numerical Harmonic Analysis. Birkhäuser, 2013.
- [11] E. J. Candès, T. Strohmer, and V. Voroninski, "Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming," *Communications on Pure and Applied Mathematics*, vol. 66, no. 8, pp. 1241–1274, 2013.
- [12] G. Colavolpe, A. Barbieri, and G. Caire, "Algorithms for iterative decoding in the presence of strong phase noise," *IEEE Journal on selected areas in communications*, vol. 23, no. 9, pp. 1748 – 1757, 2005.
- [13] A. Drémeau, C. Herzet, and L. Daudet, "Boltzmann machine and mean-field approximation for structured sparse decompositions," *IEEE Trans. On Signal Processing*, vol. 60, no. 7, pp. 3425–3438, July 2012.
- [14] M. J. Beal and Z. Ghahramani, "The variational bayesian EM algorithm for incomplete data: with application to scoring graphical model structures," *Bayesian Statistics*, vol. 7, pp. 453–463, 2003.
- [15] A. Drémeau and C. Herzet, "DOA estimation in structured phase-noisy environments: technical report," Tech. Rep., available at <http://arxiv.org:1609.03503>, September 2016.
- [16] J. M. Mendel, *Lessons in Estimation Theory for Signal Processing Communications and Control*, Prentice Hall Signal Processing Series, Englewood Cliffs, NJ, 1995.